

New coverage on digital photo editing, office productivity,
and recycling a laptop

Linux[®]

FOR

DUMMIES[®]

8th Edition

DVD includes
the full Fedora[™] Core
installation

**A Reference
for the
Rest of Us!**[®]


FREE eTips at dummies.com[®]

Dee-Ann LeBlanc

Gaming Industry Editor for
LinuxWorld magazine

Richard K. Blum





by Dee-Ann LeBlanc and Richard Blum



Wiley Publishing, Inc.

Table of Contents

Introduction..... 1

About This Book.....	1
How to Use This Book	2
Three Presumptuous Assumptions	3
How This Book Is Organized.....	4
Part I: Getting Your Feet Wet.....	4
Part II: Getting Up to Speed with Linux	4
Part III: Getting Things Done	5
Part IV: Junior Administrator Boot Camp	5
Part V: The Part of Tens.....	5
Part VI: Appendixes.....	5
Icons Used in This Book.....	6
Where to Go from Here.....	6

Part 1: Getting Your Feet Wet..... 7

Chapter 1: Getting Acquainted with Linux 9

Is Free Really Free?.....	9
Linux: Revolution or Just Another Operating System?	10
Anatomy of an Open Source Software Project	13
GNU who?	13
Who's in charge of Linux anyway?	15
Einstein was a volunteer	15
Packaging Linux: The Distribution.....	16

Chapter 2: Prepping Your Computer for Linux 19

Preparing to Use Linux and Microsoft Windows Together	20
Partitioning from scratch for a dual boot	21
Peeking at your partitions from Windows.....	22
Peeking at your partitions in Windows Vista.....	22
Peeking at your partitions in Windows 2000 and Windows XP	23
Peeking at your partitions in Windows 98	25
Making space	27
Working with Disk Partitions	29
Choosing a partitioning tool	29
Making and resizing partitions with GParted	29
Double-Checking Hardware Compatibility.....	34
Finally, Finally, Before You Get Started.....	39



Chapter 3: Installing Fedora 41

Things to Consider Before You Begin Installation	41
Getting Started with the Installation Process	42
Your First Boot	53

Chapter 4: Booting and Stopping Fedora 57

Giving Linux the Boot	58
Step 1: Power-On Self-Test (POST) leads to BIOS	58
Step 2: The BIOS passes the baton to the boot loader	59
Step 3: The boot loader (LILO or GRUB) loads the system kernel into memory	60
Step 4: Control is handed over to <i>init</i>	61
Entering Rescue Mode	62
Don't Just Turn Off the Machine!	63
Removing Linux from Your System	64

Chapter 5: Checking Out the Desktops 65

Making the Best of the Command Line	65
The GNOME Desktop	66
The menus	67
The Lock Screen tool	70
The panels	71
Changing your panels	73
Playing with GNOME desktop icons	76
Adding themes	77

Part II: Getting Up to Speed with Linux 81

Chapter 6: Working without the GUI 83

Playing the Shell Game	84
Understanding bash Command Syntax and Structure	86
Starting Programs from the Shell	87
Putting Wildcard Expansion to Good Use	88
Working with Long Commands	88
Asking Linux to complete a command or filename for you	89
Accessing your command history	90
Working with Variables	90
Variables versus environment variables	91
Checking out commonly used environment variables	91
Storing and retrieving variables' values	92
Using Redirection and Pipes	95
Redirecting command output	96
Laying pipes	97
"Help!"	97
Clearing the screen	98

Chapter 7: Getting to Know the Linux Filesystem	99
Pieces of the Puzzle	100
Touring the Linux Filesystem	100
The root of the tree	100
The importance of being/etc	102
Where temporary media lives	103
Where the computer is the /usr	103
Managing Your Filesystem without a Net (Or Mouse)	104
Viewing information about files on the command line.....	104
Understanding file listing information	107
A permissions primer	108
Comprehending file types	110
Navigating the filesystem in Fedora.....	111
Partitions versus Directories.....	113
Chapter 8: Using the Filesystem in GNOME and Nautilus	117
Sailin' through Folders	117
Setting how much you see in a folder.....	119
Creating files and folders.....	120
Opening files and running programs	121
Copying and moving files	121
Deleting files and folders.....	122
Taking out the trash	123
Viewing and changing permissions.....	123
Using CDs and Other Removable Media	124
Accessing Windows Drives on This Computer	127
Accessing Network Drives	128
Formatting Disks	129
Burning Data CDs and DVDs	129
Finding Things	131
Chapter 9: Connecting to the Internet	135
Internet Connectivity 101	135
Setting Up the Hardware	139
Selecting an Internet Service Provider (ISP)	140
Getting Information You Need from Your ISP	140
Configuring Your Connection	141
Dialin' out	142
Playing with cards	146
Look Ma, no cables!.....	148
It's All Fun and Games Until Something Doesn't Work.....	151
After You're Connected	153

Part III: Getting Things Done 155**Chapter 10: Using the Internet157**

Browse the Web with Firefox.....	157
Configuring Firefox.....	158
Them dad-gum browser plug-ins.....	163
Surfing the Web.....	164
Evolving into E-Mail	165
Setting up Evolution.....	166
Sending and checking e-mail.....	171
Taking Advantage of Instant Messaging.....	173
Using the Pidgin Instant Messenger.....	174
Troubleshooting your IM connections	176
Downloading with BitTorrent	176
Ol' Fashion File Transfers	178
Talkin' on the Phone	179
Using the default: Ekiga	180
Talking with Skype.....	181
Working with Other Internet Tools	183

Chapter 11: Putting the X in Text 185

Viewing the Contents of a Text File.....	185
Editing Text Files with nano.....	186
Saving your work as you go	186
Saving and moving on with your life.....	187
Going with gedit	188
Entering and editing text in gedit.....	188
Saving your work.....	189

Chapter 12: Word Processing and More with OpenOffice.org191

Installing the OpenOffice.org Suite	192
Word Processing with OpenOffice.Org Writer.....	193
Taking a tour of OpenOffice.org Writer	193
Working with Writer files.....	196
Spreadsheets with OpenOffice.Org Calc	197
Taking a tour of OpenOffice.org Calc.....	197
Working with Calc files	201
Presentations with OpenOffice.Org Impress	202
Using the Presentation Wizard	202
Taking a tour of OpenOffice.org Impress.....	204
Working with Impress files	208
Fine Art with OpenOffice.Org Draw	208
Taking a tour of OpenOffice.org Draw	209
Working with Draw files	212
Managing Data with OpenOffice.org Base.....	212
Getting Help from the Wizard	213
Taking a tour of OpenOffice.org base	215
Sitting down at the table	217

Layout with OpenOffice.Org Math	218
Menu bar	219
Standard toolbar	220
Document section	222
Chapter 13: Messing with Audio	225
What Sound? I Don't Hear a Thing!	225
Listening to CDs	228
Listening to Downloaded Music	230
Listening to Internet Radio	233
Ripping Music Tracks from CDs	234
Burning Data CDs and DVDs	236
Chapter 14: Messing with Video and Graphics	239
Watchin' Those Silly Web Clips	239
Checking your plug-ins	240
Adding plug-ins	241
Getting even more video support	242
Viewing Movie Files	243
Watchin' DVDs	246
Creating and Modifying Graphics	247
GIMP basic tools	248
Capturing screenshots	249
Editing an image file	250
Using GIMP scripts	252
Playing 3-D Games	253
Chapter 15: Windows-Only Media Formats and Programs	257
Commercial Software	258
CodeWeavers	258
Win4Lin	259
Cedega	259
VMware	261
Installing and Using Wine	262
<i>Part IV: Junior Administrator Boot Camp</i>	<i>265</i>
Chapter 16: Adding Software to Linux	267
Opening Downloaded Files	267
Compressing and Packaging Files to Share	269
Updating Fedora	270
Automated updates	270
Manual updates	270
Installing New Software	272
Finding More Software	276
Upgrading Your OS	277

Chapter 17: Basic System Administration279

Users and Groups.....	280
Adding new users	281
Modifying existing users.....	283
Deleting a user	284
Adding new groups	284
Modifying groups.....	285
Printing.....	286
Printer server settings.....	288
Adding a new printer	288
Modifying printer properties	291
System Monitor	295
Processes.....	296
Resources	298
File Systems.....	299

Chapter 18: A Secure Linux Box Is a Happy Linux Box301

Choosing Secure Passwords	302
Updating Software	302
Network holes	303
The Secure Shell game (SSH)	307
Software holes.....	311
Keeping an Eye on Your Log Files with the System Log Viewer.....	311
Locating Security Resources	314

Part V: The Part of Tens**315****Chapter 19: Ten Steps to Making Your
Own Wireless Access Point317**

Hello, Master!.....	318
Step 1: Plan Your Network.....	319
Step 2: Configure Your Network Card	320
Step 3: Install Your Wireless Card	321
Step 4: Install Bridging.....	325
Step 5: Build a Bridge.....	326
Step 6: Save Your Hard Work	328
Step 7: Test Connectivity.....	329
Step 8: Configure Security	330
Step 9: Configure Remote Devices	331
Step 10: Test Remote Devices	333
In Case of Difficulty	333

Chapter 20: Ten Steps to Setting Up a Samba Print Server 335

Dancing the Samba	336
Step 1: Plan Your Sharing	336
Step 2: Connect a Local Printer	337
Step 3: Manually Add a Local Printer	337
Step 4: Set the Printer Configuration	339
Step 5: Install Samba.....	341
Step 6: Configure Samba.....	341
Step 7: Start Samba	342
Step 8: Open the Firewall Port.....	344
Step 9: Let Guests Print	345
Step 10: Map a Network Printer.....	346
Troubleshooting.....	346

Chapter 21: Ten Steps to Setting Up a File Server 349

Reintroducing Samba.....	350
Step 1: Plan a Sharing Policy.....	350
Step 2: Create Linux Groups	351
Step 3: Create Folders to Share	352
Step 4: Set Advanced Folder Permissions.....	353
Step 5: Install and Start Samba.....	355
Step 6: Configure Samba.....	356
Step 7: Create UNIX User Accounts	358
Step 8: Create Samba User Accounts.....	359
Step 9: Create Shared Folders in Samba.....	361
Step 10: Map to Shared Folders.....	362
Troubleshooting.....	362

Chapter 22: Ten Troubleshooting Tips 365

Tip #1: "The Linux Installer Froze"	366
Tip #2: Checking Your Distribution Burns	367
Tip #3: "I Told the Installer to Test My Graphics, and They Failed"	368
Tip #4: "The Installer Tested My Graphics Fine, but My GUI Won't Start"	369
Tip #5: "I Think I'm in Linux, but I Don't Know What to Do!"	369
Tip #6: "I Don't Want to Boot into This!"	370
Tip #7: Changing Your Boot Environment "Permanently"	371
Changing your boot environment just for now	372
Tip #8: "I Want to Change Screen Resolutions"	372
Tip #9: "My GUI Is Hung, and I'm Stuck!"	373
Tip #10: "Help, My Machine Hangs During Boot!"	373
"Aaargh! I Forgot My Root Password! What Do I Do?"	374

Part VI: Appendixes375**Appendix A: Common Linux Commands377**

Linux Commands by Function	377
Getting Help	378
Archiving and compressing	378
Built-in bash commands	379
Files and Filesystem	380
File organization	380
File attributes	381
File locators	381
File viewers	382
File-system commands	382
mtools	383
System Control	384
Administration	384
Kernel module handling	385
Processes	385

Appendix B: About the DVD-ROM387

System Requirements	388
Using the DVD-ROM	389
Booting from the DVD-ROM	389
What You Find in Fedora 7	390
If You've Got Problems (Of the DVD-ROM Kind)	391

Index393

Chapter 1

Getting Acquainted with Linux

In This Chapter

- ▶ Diving into Linux origins and features
 - ▶ Telling Linux apart from the rest of the operating system pack
 - ▶ Depending on GNU and the GPL
 - ▶ Marveling at the Linux Company (or lack thereof)
 - ▶ Checking out popular Linux distributions
-

Ford, you're turning into a penguin. Stop it!

— Arthur Dent

Welcome to the world of Linux, the operating system developed by more than a thousand people around the world! In this chapter, you find out about Linux itself — what it is, where it comes from, and why it gets so much attention in the news these days. Prepare to have your assumptions challenged about how software *must* be developed and sold, and your mind opened to new possibilities.

Is Free Really Free?

Understanding Linux requires a radical shift of thought regarding the way that you acquire and use computer software. (**Note:** By *radical*, I mean getting to the root of the matter, rather than putting on beads and camping out in the administration building.) Your first step toward shifting your mind-set is to alter your general connotation of the word *free* to represent *freedom*, rather than *free lunch*. That's right; you can sell “free” software for a fee . . . and you're encouraged to do so, as long as you relay the same freedom to each recipient of the software.

Don't scratch your head too hard; these concepts are tough to grasp initially, especially when you consider the conditioning you've received from the commercial software industry's marketing departments. Perhaps you don't know that when you purchase most proprietary, shrink-wrapped software, you don't actually *own* the software. Rather, you're granted permission to use the software within the bounds dictated by the licensor.

Linux also has a license. However, the motives and purpose of the license are much different from those of most commercial software. Instead of using a license to restrict use of the software, the GNU General Public License (GPL) that Linux uses ensures that the software will always be open to anyone. No company can ever own or dictate the way in which you use or modify Linux — although they can have their own individual copyrights and trademarks on their various brands of it, like Red Hat and Novell. In essence, you already own Linux, and you can use it for anything you like, as long as you propagate the GPL freedoms to any further recipients of the software.

Linux: Revolution or Just Another Operating System?

Contrary to popular belief, penguins are not the salvation of modern technology. Neither do they throw parties for the urban proletariat.

— Anonymous

Author note: Cute quote . . . obviously Anonymous has never been to a Linux convention!

Before going any farther into Linux, I need to get some terminology out of the way.

Tux is the formal name of the mascot penguin that represents Linux. Rumor has it that Linux's creator, Linus Torvalds, is rather fond of these well-dressed inhabitants of the Antarctic.

An *operating system* is the software that runs your computer, handling all interactions between you and the hardware. Whether you're writing a letter, calculating a budget, or managing your recipes on your computer, the operating system provides the essential air that your computer breathes. Furthermore, an operating system isn't just one program; it consists of hundreds of smaller programs and utilities that allow us humans to use a computer to do something useful. You then run other programs (such as your word processor) on top of the operating system to get everything done.

In recent technological history, Linux has evolved from water-cooler techie chatter to a rock-solid solution for the business enterprise. The same software that was once dismissed as rogue is now being adopted and promoted by industry leaders such as IBM, Hewlett-Packard, Motorola, Microsoft, and Intel. Each of these computer manufacturers has, in some way, determined that Linux provides value for their customers (as well as for their own operations).

Linux has been accused of being “just another operating system.” On the surface, it may appear so, but if you look deeper, you can see that this isn’t so. The Linux project is a flagship leading the current trend toward open source and free (as in freedom, not free beer) software within the computing industry. A rock-solid operating system because of the model under which it was (and continues to be) developed, Linux represents much that is good in software development.

Two fundamental distinctions separate Linux from the rest of the operating-system pack:

- ✓ Linux is licensed under the unique and ingenious *GNU General Public License*, which you can read about in the next section.
- ✓ Linux is developed and maintained by a worldwide team of volunteer and paid programmers, working together over the Internet.

Linux is great for many reasons, including the fact that the folks who built it from the ground up wanted it to be

- ✓ **Multiuser:** More than one user can be logged in to a single computer at one time.
- ✓ **Multiprocess:** True *preemptive multitasking* enables the operating system core to efficiently juggle several programs running at once. This is important for providing multiple services on one computer.
- ✓ **Multiplatform:** Linux currently runs on more than 24 *platforms* (hardware types), including 32- and 64-bit Intel-based PCs, Digital/Compaq Alpha, all variants of the Apple Macintosh, Sun SPARC, the Apple iPod, and even the Microsoft XBox.
- ✓ **Interoperable:** Linux plays nice with most network protocols (languages) and operating systems, allowing you to interact with users and computers running Microsoft Windows, UNIX, Novell Netware, Apple Macintosh computers, and other, more niche groups.
- ✓ **Scalable:** As your computing needs grow, you can rely on Linux to grow with you. The same Linux operating system can run on a tiny electronic photo frame, a desktop computer, or a very large, industrial-strength server system.

- ✓ **Portable:** Linux is mostly written in the C programming language. C is a language created specifically for writing operating system–level software and can be readily *ported* (translated) to run on new computer hardware.
- ✓ **Flexible:** You can configure the Linux operating system as a network host, router, graphical workstation, office productivity PC, home entertainment computer, file server, Web server, cluster, or just about any other computing appliance that you can think of.
- ✓ **Stable:** The Linux *kernel* (the operating system) has achieved a level of maturity that makes most software developers envious. It's not uncommon to hear reports of Linux servers running for years without crashing.
- ✓ **Efficient:** The modular design of Linux enables you to include only the components needed to run your desired services. Even older Pentium computers can utilize Linux and become useful again.
- ✓ **Free!:** To most people, the most intriguing aspect of Linux is the fact that it's often available free of charge. How (the capitalists murmur) can anyone build a better mousetrap with no incentive of direct monetary return?

So where did Linux come from?

The quickest way to understand Linux is to take a peek at its rich heritage. Although programming of the Linux core started in 1991, the design concepts were based on the time-tested *UNIX* operating system.

UNIX was developed at Bell Telephone Laboratories in the late 1960s. The original architects of UNIX, working back when there were few operating systems, wanted to create one that shared data, programs, and resources both efficiently and securely — something that wasn't available then (and is still sought after now). From there, UNIX evolved into many different versions; its current family tree is so complicated that it looks like a kudzu infestation!

In 1991, Linus Torvalds was a computer science student at the University of Helsinki in Finland. He wanted an operating system that was like the UNIX system that he'd grown fond of at the

university, but both UNIX and the hardware it ran on were prohibitively expensive. A UNIX version called Minix was available for free, but it didn't quite meet his needs. So, as a computer science student, Torvalds studied Minix and then set out to write a new version himself. In his own words (recorded for posterity on the Internet because this was in an early version of an online chat room), his work was "just a hobby, won't be big and professional like GNU."

Writing an operating system is no small task. Even after six months of hard work, Torvalds had made very little progress toward the general utility of the system. He posted what he had to the Internet — and found that many people shared his interest and curiosity. Before long, some of the brightest minds around the world were contributing to Linus's project by adding enhancements or fixing bugs (errors in the code).

In this chapter, I intend to answer that last question for you. I also hope to paint a picture of the open source software development model that created Linux.

Anatomy of an Open Source Software Project

Linux isn't a product. Linux is an organic part of a software ecosystem.

— Michael Robinson, *Netrinsics*

To the casual observer (and some corporate IT decision makers), Linux appears to be a freak mutation — a rogue creature randomly generated by anarchy. How, after all, can something so complex and discipline dependent as a computer operating system be developed by a loosely knit band of volunteer computer geeks from around the world?

Just as science is constantly attempting to classify and explain everything in existence, technology commentators are still trying to understand how the open source approach can create superior software, especially in cases where there is no charge. Often the reasons have much to do with the usual human desire to fill a need with a solution. When a programmer in the Linux world wants a tool, the programmer simply writes one — or bands together with other people who want a similar package, and they write it together.

GNU who?

Imagine — software created out of need rather than projected profit. Even though UNIX ultimately became expensive proprietary software, the ideas and motives for its creation were originally based on practical needs. What people usually refer to (in the singular) as the *Linux operating system* is actually a collection of software tools that were created with the express purpose of solving specific computing problems.

The speed of Linux's popularity also wouldn't be possible without the vision of a man whom Steven Levy (author of the book *Hackers*) refers to as "The Last of the Great MIT AI-LAB Hackers" — in the original sense of the word *hacker* as someone who plays with code, not the current popular meaning that implies criminal intent. This pioneer and advocate of *freedom* software is Richard Stallman.

The Massachusetts Institute of Technology (MIT) has long held a reputation for nurturing the greatest minds in the technological disciplines. In 1984, Stallman, a gifted student and brilliant programmer at MIT, was faced with a dilemma — sell his talent to a company for a tidy sum of money or donate his gifts to the world. He did what we'd all do . . . right?

Stallman set out on a journey to create a completely free operating system that he would donate to the world. He understands — and continues to live — the original hacker ethic, which declares that information wants to be free. This concept wasn't new in his time. In the early days of the computing industry, many advancements were made by freely sharing ideas and programming code. Manufacturer-sponsored user groups brought the best minds together to solve complicated problems. This ethic, Stallman felt, was lost when companies began to hoard software as their own intellectual property with the single purpose of profit.

As you may or may not have gathered by this point, widespread and accessible source code is paramount to successful software development. *Source code* is the term for the human-readable text (as opposed to the unreadable cyber-hieroglyphics in an “executable” file) that a programmer types to communicate instructions to the computer.

Writing computer programs in binary is an extremely arduous task. Modern computer software is usually written in a human-friendly language and then *compiled*, or translated, into the computer's native instruction set. To make changes to this software, a programmer needs access to a program's source code. Most proprietary software comes only as a precompiled product; the software developer keeps the source code for those programs under lock and key.

After determining that his operating system would be built around the conceptual framework of UNIX, Stallman wanted the project name to distinguish his system from UNIX. So, he chose the recursive acronym *GNU* (pronounced *ga-new*), which means GNU's not Unix.

To finance the GNU project, Stallman organized the Free Software Foundation (FSF), which sold free (open source) software to help feed the programmers who worked on its continuing development. (Remember, we're talking *free* as in *free speech*, not *free beer*.) Although this organization and goal of creating a complete operating system was necessary and important, a much more important piece of the puzzle had to be put into place to protect this new software from big-business pirates — a concern still all too relevant today as a former Linux company tries to hijack ownership of decades of volunteer work from thousands of people around the world.

The *GNU General Public License* (GPL) is a unique and creative software license that uses copyright law to protect the freedom of the software user, which is usually the opposite of how a copyright works. Generally, a copyright is an enforceable designation of ownership and restriction from duplication by anyone but the copyright holder. When software is licensed under the GPL,

recipients are bound by copyright law to respect the freedom of anyone else to use the software in any way they choose. Software licensed with the GPL is also known as *copyleft* software (the reverse of right, get it?). Another way to remember the GPL is through its ultimate result: Guaranteed Public for Life.

While Stallman's work set the stage for Linux's rapid climb to popularity, the operating system he and his crew were working on took longer than expected. If you're interested in the completed version, go to www.gnu.org/software/hurd/hurd.html.

Who's in charge of Linux anyway?

As an open source project evolves, various people emerge as leaders. This leader is often known as the project's *benevolent dictator*. A person who becomes benevolent dictator has probably spent more time than anyone else on a particular problem and often has some unique insight. Normally, the words *democratic* and *dictator* are never paired in the same sentence, but the open source model is a very democratic process that endorses the reign of a benevolent dictator.

Linus Torvalds is still considered the benevolent dictator of the Linux *kernel* (the operating system's core). He ultimately determines what features are added to the kernel and what features aren't. The community trusts his vision and discretion. In the event that he loses interest in the project, or the community decides that he has gone senile, a new leader will emerge from amongst the very competent people working with him.

Einstein was a volunteer

Someone who is a volunteer or donates time to a project isn't necessarily providing a second-rate effort (or only working on weekends and holidays). In fact, any human resources expert will tell you that people who choose to do a job of their own free will produce the highest quality products.

The volunteers who contribute to open source projects are often leaders in their fields who depend on community collaboration to get useful work done. The open source concept is no stranger to the scientific community. The impartial peer-review process that open source projects foster is critical in validating some new feature or capability as being technically correct.



Those who paint the open source community as copyright violators and thieves often misunderstand — or outright ignore — these vital issues. Open source programmers are very proud of their work *and* are also very concerned about their own copyrights, not wanting their work to be stolen by others — hence licenses such as the GPL. This concern creates an atmosphere with the

greatest respect for copyright. Bandits who claim that they're "just being open source" when they steal other people's hard work are grossly misusing the term to soothe their own consciences.

Many have also pointed out that if copyright is violated in open source, it's easy to tell. Watch the news and notice how often large software corporations are convicted of stealing other people's code and incorporating it into their own work. If the final product is open source, it's easy for anyone to look and make sure nothing stolen is in it. As you might imagine, tracking down such copyright violations is much more difficult in a closed source scheme.

Packaging Linux: The Distribution

What people call a *Linux distribution* is actually the culmination of the GNU project's tools, the Linux kernel, and any number of other open source (and closed source) software projects that sprang up along the way.

Robert Young, cofounder and current chairman of Red Hat, has coined an analogy comparing Linux to ketchup. Essentially, the operating system called Linux — including the GNU tools, Linux kernel, and other software — is a freely available commodity that, like ketchup, different distributors can package and label in different containers. Anyone is encouraged to package and market the stuff, even though the ingredients are fundamentally the same.

Linux is a complex, malleable operating system, and thus it can take on many appearances. The greatest Linux advancement in recent years has been easier installation. After all, the tools that today enable the casual PC user to install Linux weren't originally available. Companies such as Red Hat saw this as an opportunity to add value to an existing product, and the concept took off like gangbusters.

To draw again on the ketchup analogy, various distributions of Linux have a slightly different *flavor* or texture; your distribution preference may be spicy, mild, thick and gooey, or runny. However, you can rest assured that any of the following distributions have the same Linux and GNU heart and soul. Each short description in this list includes a Web address where you can find more information about each project:

✓ **Debian GNU/Linux:** This distribution — one of the oldest — is a recognized favorite among advanced technical circles. Historically, it's relatively difficult to install. Easier to install (and use) distributions with Debian underneath are Knoppix, Ubuntu, Xandros, and Linspire (formerly known as Lindows). Of these, Ubuntu is by far the most popular, with a dedicated worldwide fan base. Debian, Knoppix and Ubuntu are free as in free beer, while Xandros and Linspire are both commercial

distributions — though each has a free version you can try — created for those who really are most comfortable in Windows and don't want to have to master the nitty-gritty of using Linux.

Knoppix is particularly interesting as it's a live distribution, meaning that you take a Knoppix CD-ROM or DVD-ROM, place it in the drive, and reboot the machine into a fully operational Linux desktop. This book contains another live distribution option for those who don't want to install Linux on their hard drive. See Chapter 3 for more information. Knoppix, Linspire, Ubuntu, and Xandros are covered on the DVD-ROM that comes with this book.

```
www.debian.org
www.xandros.com
www.ubuntulinux.org
www.knoppix.org
www.linspire.com
```

- ✓ **Mandriva:** This distribution demonstrates the power of the GPL by allowing this competing company to stand on the shoulders of giants. Mandriva (formerly known as Mandrake) was originally based on Red Hat Linux (something that simply could not happen in a closed source environment), but has since become an excellent solution in its own right. Mandriva offers a wide range of commercial products and services as well as its free version.

```
www.mandrivalinux.com
```

- ✓ **Red Hat and Fedora:** Red Hat claims the prize for successfully mass marketing the Linux operating system. Red Hat has validated Linux by packaging the GNU and Linux tools in a familiar method of distribution (shrink-wrapped) and has included value-added features to its product, such as telephone support, training, and consulting services. The commercial version of this distribution is Red Hat Enterprise Linux, and the free version is called Fedora. Fedora is the distribution covered in this book.

```
www.redhat.com
www.fedoraproject.org
```

- ✓ **Slackware:** Of all the more widely recognized surviving Linux distributions, Slackware has been around the longest (in fact, its installation interface remained unchanged until about a year ago). Slackware has a very loyal following, but isn't well known outside of the Linux community. Like Debian in terms of spirit, the Slackware crowd is as respected in Linux circles as the weathered old-timers who share stories of carrying around a shoebox full of diskettes. This distribution is also free.

```
www.slackware.com
```

- ✓ **Novell and SuSE:** (Pronounced *soo-za*) This distribution originally hails from Germany, where it has a very loyal following, and was purchased by Novell, Inc., as part of the company's new focus on Linux. Novell

offers both a commercial distribution in SuSE Enterprise Linux, and a free distribution with openSuSE, an effort similar effort to that of Fedora.

www.opensuse.org
www.novell.com/linux/

As you can see, many paths (in the form of distributions) lead to Linux. It's important to note that regardless of which distribution you choose, you're using the same basic ingredients: the GNU tools and the Linux kernel. The major differences you'll encounter among distributions are

- ✓ **Installation programs:** Each distribution has developed its own installation program to help you achieve a running computer system. Some installation programs are designed for the casual computer user (hiding the technical details); others are designed with the seasoned system administrator in mind. Some of the simpler ones offer an “expert mode” for those who want to have more control right from the beginning.
- ✓ **Software versions:** Different distributions may use different versions of the *kernel* (the core of the operating system) and other supporting software packages — which makes for a plethora of versions. Open source projects are dynamic and release new versions regularly, as opposed to the often-sluggish development cycle of traditional commercial software.
- ✓ **Package managers:** Even though one Linux program should be able to run on any distribution, tools called *package managers* keep track of the software on your system and ensure that you have all the required supporting software as well. Distributions are usually dependent on one particular package manager. More recent in the grand scheme of things, package management has also come to involve adding easy-to-use update routines, as well as an easy way to add software without the need of a computer science degree. Chapter 16 provides more information about package management.

It would be impossible to account for *every* possible installation of every Linux distribution. Okay, maybe not literally impossible, but you'd need a forklift to bring your *Linux For Dummies* book home from the bookstore if I did. Consequently, I try to summarize the concepts needed to install any Linux distribution into this one book with enough detail to get you through the process. (As you can imagine, that's a bit of a challenge!)



I chose Fedora as the sample distribution because Red Hat has become a recognized Linux standard, and its Fedora project is specifically aimed at home and small business users who cannot afford (or have no need) to purchase higher-level products. Even better, if you do use Red Hat Enterprise Linux (RHEL) in your office or organization, Fedora is a proving ground for the technologies that will make their way into RHEL.



Check out Appendix B for a full list of the goodies you can find on the DVD.

Chapter 2

Prepping Your Computer for Linux

In This Chapter

- ▶ Taking basic preinstallation steps
 - ▶ Using Linux and Windows on the same computer
 - ▶ Customizing disk partitions before installation
 - ▶ Knowing (and finding) your hardware information
 - ▶ Preparing for DVD or CD installation
-

You got to be careful if you don't know where you are going, because you might not get there.

— Yogi Berra

Because there are three major approaches for installing Linux, look at this chapter introduction as a quick road map that tells you which sections are important for your particular situation. If you have a spare computer that's only going to run Linux and nothing else, you're in luck! You can skip the "Preparing to Use Linux and Microsoft Windows Together" section. In fact, if you're feeling brave, you may want to skip right to Chapter 3 and start your installation. There's troubleshooting information in Chapter 22 as well.

Of course, many people don't have the luxury of having more than one computer. There are many ways to approach a *dual booting* situation, where you have both Microsoft Windows and Linux installed on the same computer and reboot the machine to switch between one and the other. The scenario that many people prefer is to add a second, fresh hard drive to install Linux on — or to remove all the data from an existing drive and erase it for Linux. If you decide to go with one of these options, make a note of which drive is which as far as the computer sees them: You want to make sure that you leave your Microsoft Windows installation untouched. All you need to know is which drive (Windows or Linux) is first and second as far as the computer is concerned. You can find out about how to read the order information in the section "Peeking at your partitions from Windows" later in this chapter. When you are sure you know which drive is which, proceed to the section "Double-Checking Hardware Compatibility" later in this chapter.



If you are removing information from a hard drive so you can erase it, be sure that you have moved everything before installing Linux, because Linux is going to completely wipe the drive.

Those who can't spare an entire hard drive for Linux and already have Microsoft Windows installed will have to resize their current Windows installation. If this is your situation, you need to work through this entire chapter. Absolutely **DO NOT PROCEED TO CHAPTER 3** without at least reading the next section, "Preparing to Use Linux and Microsoft Windows Together." Sorry for yelling, but you can wipe out your whole Windows installation if you don't take precautions!



If you are using Microsoft Windows Vista and plan to purchase a product to help you do things like resize partitions, make sure and read the product's box and see if Vista is supported. Partition Magic, for example, has an upgrade you can download to add Vista support. See www.symantec.com/windowsvistasupport.

If you absolutely don't want to dual boot, you have three other options — I know I said that there were four approaches in total and adding three here brings us up to six, so give me a moment to explain. You can use VMware (see Chapter 15) to install a "virtual" Linux machine that lives in a window inside your Windows installation. You can install only Linux on the computer and then use VMware or Win4Lin (see Chapter 15) to install a virtual Windows machine that lives in a window within your Linux installation. Or, if the thought of changing anything on your computer gives you hives, you can use a *live distribution* (see Chapter 3) to boot your computer into Linux without installing anything.

So give some thought to which of the options detailed here you prefer, and then read on.

Preparing to Use Linux and Microsoft Windows Together

If you're planning to run Linux and Microsoft Windows on the same machine, the odds are that you already have Windows installed and have been using it for some time. Because I hate to hear screams of anguish all the way up here in western Canada, take a moment to assess what you have and what you need to do.



On the off-chance that you actually don't have Windows installed yet and still want that dual-boot capability, you should install Windows *before* you install Linux. Otherwise, during installation, Windows overwrites the part of your hard drive that Linux uses to store its *boot menu*. (This factor can create a mess later when you want to boot back into Linux!) Those who need to install Windows should skip down to the section “Working with Disk Partitions” to find out how to set up your Windows installation so it causes the least fuss when it's time to add Linux. Then after you have Windows installed, return here.

The majority of you, however, want to dual-boot because you've got one machine and it's already running a Windows installation that you really don't want to redo. If you intend to put Linux on a brand-new hard drive that has nothing on it already (or that you are 100 percent confident that you can erase), skip down to the section “Working with Disk Partitions.” However, if you need Linux and Windows to share the same hard drive and you already have Windows installed, you have some extra work to do in order to make sure that everything goes smoothly. The rest of this chapter focuses on getting you through this process.

Partitioning from scratch for a dual boot

Again, if you plan to take a fresh hard drive (or erase an existing one) and install both Windows and Linux on it, *be sure to install Windows first*. While you're going through the Windows installation, you are asked to *partition* your drives, which refers to using software to break them into smaller virtual drives. There are three types of partitions available: primary, extended, and logical. A hard drive can have three *primary* partitions and one *extended* partition. Inside that extended partition, you can have up to 12 *logical* partitions — think of an extended partition as just a cardboard box that contains the logical partitions. Logical partitions hold data; extended partitions just hold logical partitions. Because I can't predict what software you want to install, I recommend having at least 10GB of space available for your Linux installation. More is always better because it gives you more room for downloads and even more programs, and, at the very least, you can install everything from this book's companion DVD-ROM into that 10GB.



Make a note of the partition you dedicate to Windows and the one you dedicate to Linux. In the note, jot down which hard drive each partition is on (the first, second, third, and so on) and the partition number each is on the disk (again, first, second, and so on). You'll need this information when installing Linux.

Those who aren't starting from scratch for a dual boot likely need to make changes to their current installation. Proceed to the next section to find out how.



You are not limited to a dual boot. You can have three or more operating systems on the computer if you want — and if you have room.

Peeking at your partitions from Windows

Before you change anything, you must collect some important information about your current setup. You need to know two major things about your hard drives before proceeding. The first is to know whether you have any unpartitioned space left on your hard drive(s) — and, if so, how much space you have. If you discover at least 10GB of unpartitioned space, you can skip straight to the “Double-Checking Hardware Compatibility” section later in this chapter.

If you find using the tools in this chapter awkward, you can purchase a commercial program called Partition Magic (www.powerquest.com) to make the process easier. I’ve chosen to cover a free alternative in this chapter instead.

Peeking at your partitions in Windows Vista

The Vista operating system has security more on its mind than previous versions of Microsoft Windows did. As such, to use the instructions that follow, you need to be using an account with administrative privileges. After you have logged into such an account, do the following:

- 1. Open the Control Panel by choosing Start→Control Panel.**

The Vista Control Panel opens.

- 2. Select Classic View if necessary.**

- 3. Double-click Administrative Tools.**

The window displays the various administrative tools available in Vista.

- 4. Double-click Computer Management.**

The window displays the various tools available in Vista to manage your computer. You may need to click Continue to proceed to this dialog box.

- 5. Within the left pane, select Disk Management.**

The Disk Management tool opens (Figure 2-1).

Figure 2-1:
The
Windows
Vista Disk
Manage-
ment tool.



If you look through the Disk 0, Disk 1, and Disk 2 listings, you see that one of the boxes says Vista (C:). That's the Vista installation. You also see some boxes that say Unallocated. They are free space that isn't being used for anything. The EFI items you see are for an Apple Intel Macintosh installation.

Those who find an unallocated partition that is 10GB or larger (refer to Figure 2-1), make a note of which disk this partition is on and what numbered partition it is on that disk, and skip ahead to the section "Working with Disk Partitions." If you don't have any (or enough) unallocated space, keep this window open and jump to the section "Making space" later in this chapter.

Peeking at your partitions in Windows 2000 and Windows XP

The Windows 2000 and XP operating systems only allow certain accounts to control and secure the files and folders; to get information about the computer's disk-space usage, first you have to be logged on as the Administrator (or use an account that has administrative permissions). After you log on, you use the Computer Management application in the Administrative Tools collection to find the details of the computer's hard-drive usage. Follow these steps:

1. Open the Control Panel by choosing Start→Settings→Control Panel.
2. Open the Administrative Tools folder and double-click the Computer Management icon.
3. In the left pane of the Computer Management application, click the Disk Management folder icon.

Within a few seconds, the right pane displays the current status of the storage devices on your computer, such as the hard drive(s), CD drive(s), DVD/CD drive(s), and so on. Figure 2-2 is an example of the Windows XP Disk Management display, and Windows 2000 users see a nearly identical view.

Figure 2-2:
Disk
Manage-
ment
information
within the
Windows XP
Computer
Manage-
ment tool.

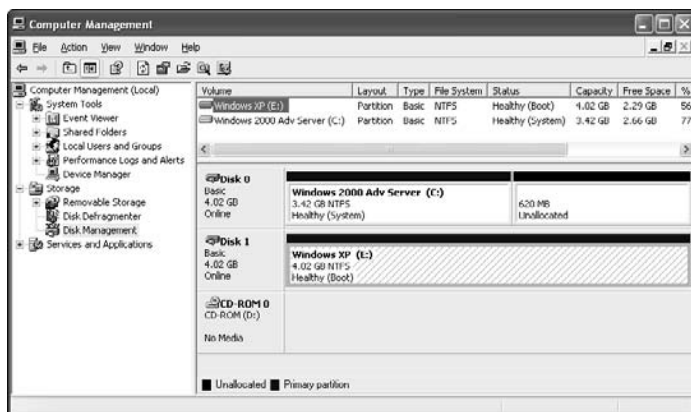


Figure 2-2 reflects a computer that has two hard drives and one CD drive. The important thing to look for here is the word *Unallocated* in the Disk listing — an example is shown in this figure. Unallocated partitions are not assigned to any operating system and are free to use for your Linux installation.

If you do find an unallocated partition (refer to Figure 2-2), you're all set if it's 10GB or larger. Make a note of which disk this partition is on and what numbered partition it is on that disk, and skip ahead to the section "Working with Disk Partitions." For example, in Figure 2-2, the partition is on disk 1, and it's partition 2.

If you don't have any (or enough) unallocated space, keep this window open and jump to the section "Making space," later in this chapter.

Peeking at your partitions in Windows 98

Windows 98 unfortunately does not provide a graphical tool like Computer Management for Windows 2000 and XP. Instead it uses `FDISK`, a command-line tool that indicates the partitions on your hard drive. To find out the details about a computer's hard drive in Windows 98, follow these steps.



Be careful in `FDISK`: You can wipe out your data if you make any changes with it. When you are exiting the program, be sure not to save any changes.

1. **Open an MS-DOS prompt window by choosing Start⇨Programs⇨MS-DOS Prompt.**

2. **Type `FDISK` and press Enter.**

In Windows 98, you can enter `FDISK` in uppercase, lowercase, or any mixed case you like as long as you spell `FDISK` correctly!

You will very likely be prompted to display large “disk” (drive) information — “large” is relative to what was a big hard drive in the time of Windows 98, which means a drive larger than 512MB. If you don't see the prompt, skip to Step 4, and the `FDISK` menu options will be similar to Figure 2-3.

3. **If you see the large disk prompt, choose Y at this prompt and then press Enter.**

The `FDISK` menu options appear.



Figure 2-3:
FDISK
menu
options.

4. **Display the current drive-partition information.**

If you have more than one hard drive in your computer, `FDISK` displays a fifth menu choice so you can change between disks.



Here's how to use the fifth menu choice to change to another disk:

a. Type 5 and press Enter.

The FDISK screen displays all the hard drives in your system.

b. Type the number of the disk you want and then press Enter.

The top of the FDISK menu screen displays the number of the drive that FDISK is working with.

5. To display partition information for the disk number displayed, type 4 and press Enter.

Figure 2-4 shows an example of the Display Partition Information screen within FDISK.




Figure 2-4:
Partition
information
in FDISK.

Figure 2-4 indicates that no free, unallocated disk space is on the drive. If the disk had unpartitioned space, FDISK would display the amount of space available.

However, in Figure 2-4, you see that a portion of the disk space is configured as an extended partition EXT DOS (see the section “Partitioning from scratch for a dual boot” earlier in this chapter for a breakdown of the partition types). Unfortunately, FDISK doesn’t tell you anything about what’s inside this extended partition within this screen. You have to dig deeper.

6. To view the logical partitions within the extended partition, select the EXT DOS entry and press Enter.

Figure 2-5 is an example of a logical-partition setup. In this case, all space within the extended partition is currently assigned. You can tell whether there’s free space by totaling the Usage percentages for the drive. That total determines what you do next:

- If the total is 100 percent, no space is available. If this is the case (or if you have less than 10GB available), leave this tool open and go to the next section, “Making space.”
- If you have 10GB or more available, skip to the section “Working with Disk Partitions.”




Figure 2-5:
Logical
drive
information
in FDISK.

Making space

You may need to ultimately resize your partitions for a number of reasons. You may have some unallocated space, but less than 10GB of it. You may have 10GB of unallocated space, but it's broken up into numerous pieces across your drives and needs to be consolidated. Or (most commonly), you may not have any unallocated space at all. In all these situations, you need to move data around so that you have room to install Linux and store any files you want to keep in your Linux filesystem.



You probably want *more* than 10GB of space. For example, if you download lots of multimedia, you'll quickly eat up whatever was left after you installed your software! So, 20GB might be a better amount to shoot for — *at minimum*.

The key to making your space is (first) to determine how much of your drive space you're really using and jot this information down for use later.

- ✓ Remember, do not try to resize Windows Vista partitions unless a tool specifically states that it can safely do so.
- ✓ In Windows XP and 2000, return to the Computer Management tool discussed in the earlier section “Peeking at your partitions in Windows 2000 and Windows XP” and look in the Capacity, Free Space, and %


columns. Here you can see how much space you have already used in each partition, and how much you have available.

- ✓ If you're bringing a venerable Windows 98 and 95 machine back to dual-boot life, use the Computer Management tool this way:

1. **Double-click My Computer to open the Computer Management tool.**
2. **Choose View→Web Page.**
3. **Highlight your first drive.**

You see something like what is shown in Figure 2-6.

Figure 2-6:
How much
disk space
is available
on this
Windows 98
user's
C: drive.



Determining how much space is enough for your Windows setup — you don't want to shrink it so much that you run out of space in Windows — also involves taking a look at how you use your machine. Again, if you download and save a lot of multimedia stuff, it's wise to keep enough room around to accommodate your downloading habit.

If you find that your Windows drives are all too full and you have no room to spare (or you download and save too many audio, video, or graphics files), it may be time to burn some of your files off onto CDs or DVDs so you can delete them from your hard drive. The other option is to add a second drive to the computer — if you're not comfortable doing this yourself, many stores will install a drive for you for a fee.

When you've decided which partition(s) you want to resize, make a note of which hard drive each is on and which partition each is on that hard drive.

Working with Disk Partitions

In the section “Preparing to Use Linux and Microsoft Windows Together” earlier in this chapter, I discuss why you might need to set aside disk space for Linux on a computer that’s already running Windows. The techniques I cover here assume you’ve found free space that you can give to Linux after you resize your partitions. If you’re installing on a fresh, additional hard drive, are wiping out a whole hard drive to give it over to Linux, or already have a chunk of unallocated/free space bigger than 10GB, skip ahead to the section “Double-Checking Hardware Compatibility,” later in this chapter.



If the thought of changing anything on your Windows drive makes you queasy, a cheap, additional hard drive can be worth its weight in gold.

Choosing a partitioning tool

Perhaps the easiest way to adjust your existing partitions is through the use of a commercial program such as Partition Magic. Partition Magic enables you to view the partition information for your hard drive(s). You can also use this utility to resize, move, and add partitions. If all you want to do is resize, however, you don’t have to go out and buy software just to get that done. Instead, you can download a free tool to do the job for you (see the next section for how to do so).



No tool is perfect. Before going any farther, make a backup of anything you don’t want to lose off your Windows installation! Things can go wrong during the various steps of partitioning, no matter what tools you choose.

Making and resizing partitions with GParted

If you found that you have free space but that you need to consolidate it from a number of smaller pieces into a bigger one, you need to resize some of your partitions. The program you use for this purpose is GParted. This utility can work with Windows partitions, including at least the original release of Microsoft Vista.

Getting gparted

You need to download the `gparted` program before you can use it, and burn it onto a CD-ROM. To download `gparted`, do the following:

1. Go to <http://gparted.sourceforge.net/>.

2. Click the LiveCD link on the left.

You're taken to the `gparted` LiveCD page, where you can download an ISO image (think of it as a zip file that contains everything needed for the CD) for a CD-ROM that you can boot a computer with. After you boot, you find yourself in a Linux session where you can use `gparted` to adjust your partitions.

3. Look for the sentence, Download the latest version here, and click the link within.

At the time of this writing, you click the word [here](#). Clicking takes you to a page listing all the versions of the LiveCD.

4. Under the latest (top) version, click the file ending in .iso.

It will look something like `gparted-livecd-0.3.4-5.iso`. If you only see a list of names with numbers and no `.iso` at the end, click the top one, and the folder will open to show you the file. When you click the file, your browser may or may not ask what to do with it. Choose to download the file. Compared to most ISO images, this LiveCD image is small, only 52.8MB. It's designed to fit on business card-sized CDs.

Burning your ISO onto a CD-ROM

You need to have a CD-Writer, or a CD-Rewriter to continue (or talk to a friend who has one). You also need to have writeable or rewriteable CD-ROMs. Look for the symbols CD-W (writer) or CD-RW — you can also use a DVD burner and DVDs if you need to, look for the symbols DVD-W or DVD-RW.

When you have the file downloaded, you'll need software you can burn it with. Most DVD burners these days come with a program called Nero Burning ROM (www.nero.com). If you don't already have it and know how to use it, you can download a trial version of Nero 7. Then go to the Nero site and click Downloads↔Trial Version↔Nero 7. There you can find the User Guides if you need them.



The handy part of using Nero in Windows is that you can click the Downloads link on this site and find a Linux version of this same tool. Be sure to download the RPM version and read the QuickStart guide if you do so.

***Altering your partitions with gparted***

Before proceeding, double-check that you have backed up all of your important data. There is always a risk when altering things at this level on a hard drive that something may go wrong.

If you made any notes about what you want to change as you prepared to adjust your partitions, pull them out now. Then do the following:

1. Put the **gparted** CD into the CD- or DVD-ROM drive and reboot your computer.

A menu appears, offering you different ways to boot the computer.

2. Choose the top option and press **Enter**.

Text scrolls by as the LiveCD figures out some of the hardware on your system. Then the scrolling will pause with the text `>> Loading keymaps` and offer you options.

3. Type in the number for the appropriate option if you don't want U.S. English, and press **Enter**. If you do want U.S. English, just press **Enter**.

The system notes your selection and then allows you to choose a language.

4. Press **Enter** for U.S. English, or type in the number for the language you want and press **Enter**.

More text scrolls by. When the LiveCD finishes loading, you immediately see the GParted tool displaying the contents of your first hard drive (see Figure 2-7).




Figure 2-7:
The main
GParted
window.

5. On the right, you will see a large button with a hard drive and its size. Click that button and select the drive you want to work with from the drop-down list.

If necessary, refer to your notes (see the section “Preparing to Use Linux and Microsoft Windows Together,” earlier in this chapter) to see whether it’s the first, second, third, or so on of your hard drives. If you have both IDE and SCSI hard drives, you can tell the difference by the “hd” starting IDE names and the “sd” starting SCSI names. (Most home users have IDE

drives; you usually have to go out of your way to get SCSI, which are more expensive.) The order they're installed on the machine is indicated alphabetically. For example, hda is the first IDE hard drive, hdb is the second, and so on. (It's the same for SCSI, sda and sdb.)


After you make your selection, GParted displays the chosen drive's contents.

6. In the lower section of the window, select the partition you want to **resize**.

Again, refer to your notes from the section "Preparing to Use Linux and Microsoft Windows Together," earlier in this chapter, to see which partition you want to change.

7. Click **Resize/Move** to open the **Resize/Move** dialog box (see Figure 2-8).

Figure 2-8:
GParted's
Resize/
Move
dialog box.




8. **Make your sizing changes.**

Note that you have multiple options here. You can choose Free Space Preceding (MiB) to resize the partition and leave the specified amount of empty space in front of it, New Size (MiB) to specify an exact new size for the partition, or Free Space Following (MiB) to resize the partition so that there is the specified amount of empty space after it. If the Resize/Move button doesn't become available, click in another text box within the window.



If you have smaller pieces of blank space scattered through your drive between partitions, start by absorbing them into the other partitions. Do so by changing the Free Space Preceding (MiB) or Following (MiB) entry for each partition to zero, thus squeezing things together more efficiently (by moving all your free space to the end of the drive). This concept is perhaps better illustrated with an image, so see Figure 2-9.

Figure 2-9:
Combining
2GB and
3GB free-
space
partitions by
absorbing
the free
space and
then
creating
new free
space to
make one
10GB
partition.



9. Click Resize/Move to save the changes for that partition.

At the bottom of the screen, the change you just requested appears in a list of pending changes. Your hard drive isn't permanently altered yet.

10. Return to Step 5 if you have more drives or partitions you want to change. Otherwise, proceed to Step 13.

11. Click Apply to put the changes into effect.

The Apply Operations to Hard Disk dialog appears.

12. Click Apply to confirm that you want to make your changes.

Now is when your hard drive is actually changed! The Applying Pending Operations dialog box appears, showing you what is happening as the changes are made.

13. When the changes are complete, click Close.

The dialog box closes, as does the pending changes list at the bottom of the screen.

14. Double-check that you have given yourself at least 10GB of free space in total, if not more.

If not, resize things further.

15. Reboot the machine, removing the GParted CD or DVD in the process.

If something went wrong and you're having a hard time booting the system, use your Windows system-recovery disc (such as your XP or Vista CD) to repair the damage. Fortunately, such damage is very rare!

After you've finished your repartitioning, proceed to the next section to continue gathering the remaining information needed for the installation.

Double-Checking Hardware Compatibility

If you're installing Linux on hardware you already own, just give it a try and then see what doesn't work. Parts of this section address fixing hardware problems in a general sense. In other chapters, more specific items such as sound cards (Chapter 13), wireless cards (Chapter 9), and so on, are addressed. So, if you run into trouble, start in the specific sections dedicated to particular tasks, and then come here for more general help if you still haven't solved the problem.

The biggest problem areas are wireless cards (which the Fedora project is heavily targeting for Fedora 7) and the *very* latest whiz-bang multimedia hardware such as video cards, sound cards, and image scanners. You can check hardware compatibility lists before installing or purchasing new hardware, but they are of limited use because the hardware world changes so quickly. If you're interested in looking, use the Red Hat Enterprise Linux listings at <http://hardware.redhat.com/hcl/>. (There's no official Fedora list.) Keep in mind that this list focuses on business equipment instead of home items, so just because you don't see something listed there doesn't mean it's not supported.



Don't worry about items being *Certified* (heavily tested to make sure they work properly) or not. *Supported* and *Compatible* are fine most of the time for a home user. Ultimately, the best way to tell if a piece of hardware is supported or not is to do a Web search. Many go to www.google.com/linux and do a Web search on the make and model of the hardware, plus the word **Linux**. For example, you can search on Innovision DX700T Linux to find out how other people are faring with this particular brand and model of hardware. (No endorsement of any hardware mentioned as examples is implied here, of course.) Such a search is likely to show you what problems and successes people have run into with that particular piece of hardware.



If even thinking about computer hardware gives you dizzy spells, don't worry, you can find plenty of information about on the Internet. A great place to start is www.tomshardware.com. Other places to look for information on how various devices work in Linux include

Laptop considerations

The current distributions of Linux do very well on relatively new notebooks and laptops. (See www.linux-laptop.net for an excellent research site on how Linux gets along with various makes and models.) If your laptop is a common brand, you shouldn't encounter any problems installing Linux. However, laptops often contain WinModems. (Hardware labeled with the *Win* prefix is only for Windows, and can't figure out what to do with Linux.)

If you plan to purchase a laptop for Linux, check out its modem and other hardware (such

as network cards) to make sure they're not Win-branded. If the built-in or default hardware for the laptop is Win-labeled (or you discover while researching the machine that it contains a Win product, even if it isn't properly labeled), you might be able to switch the offending hardware for a PC (or PCMCIA) Card. Most current laptops contain at least one PC Card slot to give you a place to slip in a PC Card modem, network card, or combo modem-network card. As long as you stick with a common brand of PC Card, it should work well with Linux.



✓ **Other Linux-oriented Web sites:** In particular, the generic Linux hardware list at www.tldp.org/HOWTO/Hardware-HOWTO/.

✓ **Vendor Web sites:** Many hardware vendors do support Linux, but they don't make it easy to find information about it. In general, go to the vendor's forums for the piece of hardware if they have them, the FAQ for the hardware, or follow the Support links to look for downloads for Linux. Don't download what you find if there are downloads. The point is to look and see if they exist. The *driver* (the software that tells the operating system how to use the hardware) available for download might actually be included in your Linux installation. Only download the driver from the vendor if this is the only way you can get it.

If worst comes to worst, you might not find any information on the hardware in question involving Linux. However, again, this doesn't mean that the hardware won't work. Try it anyway if you already have the item. You find that it works fine. Or, you might not be able to use the very latest features, whereas the rest works just fine (for example, with a latest-generation video card, the newest fancy features might not work, but you can still use it as a generic SVGA at the very least).

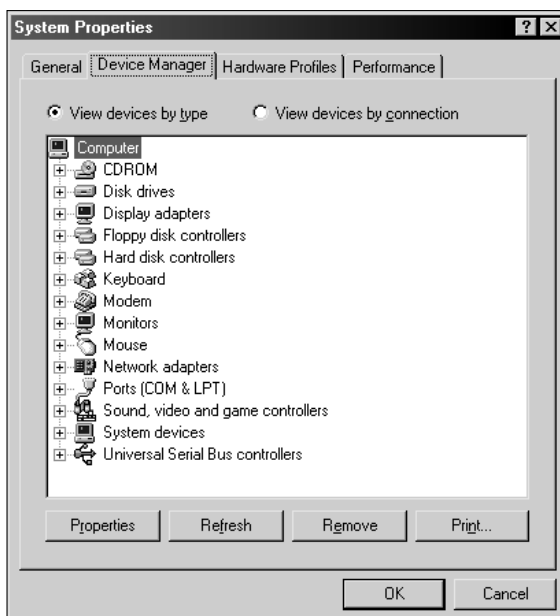
✓ **The dreaded manuals:** When possible, keep your computer manuals (especially those for your video card and monitor) handy, just in case you need them in order to answer a question asked by the installer — most people don't have to deal with this at all, but some do.

If you need to find out exactly what hardware is in your machine, you have the following options:

✔ **Use an existing operating system to document your hardware.** If your computer is already running Windows, you can collect a lot of information from the Windows environment. Use one of the following methods, depending on your system:

- **In Windows 98:** (Yes, some folks are still running it.) Choose Start⇨Settings⇨Control Panel⇨System⇨Device Manager to access the dialog box shown in Figure 2-10.

Figure 2-10:
In Windows
98, the
Device
Manager
gives you
information
on what
hardware
you have
installed.



- **In Windows 2000:** Choose Start⇨Control Panel⇨System⇨Hardware⇨Device Manager to access the list of hardware installed on your machine, as shown in Figure 2-11.
- **In Windows XP:** Choose Start⇨Control Panel⇨Printers and Other Hardware. Here you can select one of the items from the dialog box under Pick a Control Panel Icon, or you can look to the See Also section and select System. If you choose System, from there you choose Hardware⇨Device Manager to access the list of hardware installed on your machine (as shown in Figure 2-12).

You can double-click each item within the Device Manager to display the corresponding details.



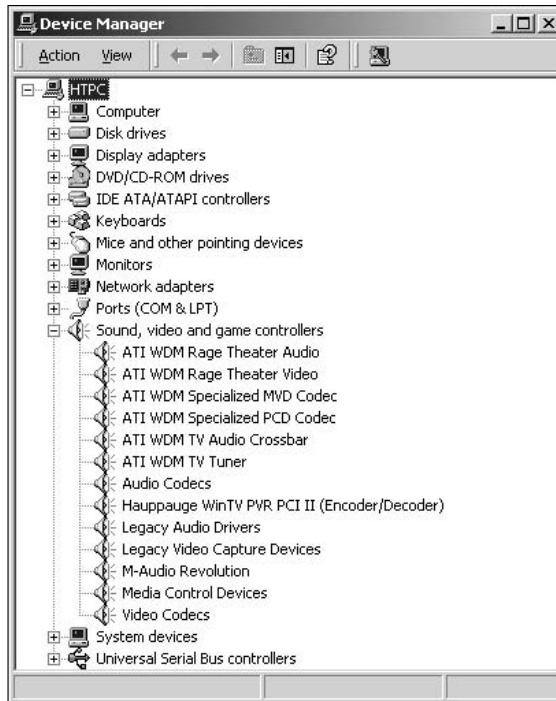


Figure 2-11:
The
Windows
2000 Device
Manager.




Figure 2-12:
The
Windows
XP Device
Manager.

- **In Windows Vista:** Choose Start⇨Control Panel⇨Device Manager to browse your hardware. A warning dialog box will appear; click Continue.

✓ **Download a PC hardware-detect tool.** If you don't have any diagnostic tools and you have a relatively current version of DOS, you can download various PC hardware-detect tools, such as *Dr. Hardware*, from the Internet. The *Dr. Hardware* tool contains lots of information about what's inside your machine. This tool is shareware, and the usage and fee information is available from the Gebhard Software Web site (www.dr-hardware.com/).

✓ **Gather information by reading the screen when the computer starts.** If your system doesn't contain any operating systems and you don't have any of your system's documentation, you can resort to reading the screen as your computer starts. On some systems, the video information is displayed as the computer boots. You may have to reboot several times to read the information if it goes by too fast. Also, some systems display the PCI components and their settings as the system is starting up. Again, you may have to reboot several times to gather all the information.



You can try pressing the Pause-Break key on your computer (it should be near the Scroll Lock key) to freeze it during boot. You can then unfreeze it by pressing any key when you have finished reading.

✓ **Access the Basic Input/Output System (BIOS) information.** Stored in a small area of memory and retained by a battery, this is sometimes referred to as CMOS (Complementary Metal-Oxide Semiconductor), which indicates the type of computer chip that can store and retain information. The amount of information stored in the BIOS can range from very little to quite a lot. Some newer systems may display several screens of information about the computer's hardware.

If you choose to access the BIOS, make sure you do so before *any* operating systems load. Most manufacturers indicate the keyboard key (or key sequence) that gets you into the BIOS (or Setup) on-screen when the system is starting up — for example, *Press Del to enter Setup*. If you can't find the keyboard sequence, check the manufacturer's Web site. After you've entered the BIOS, you typically navigate around with the arrow keys, Tab key, or Enter key. Some BIOS environments also use the function keys; look for a list of function-key options at the top or bottom of the screen.



Be especially leery of labels on hardware boxes and Web sites that include the term *Win* (as in *Windows*). These components, such as *WinModems*, rely on Microsoft Windows to be able to function — even worse, the packaging may show nothing that suggests this limitation. Only a very slight chance exists that you can find a Linux driver for *Win* hardware. If you do find one, copy it to a floppy *before* you install Linux. If you can't find a driver and you

need to use a modem, put down a little cash and get a modem that is supported properly. (For more information about WinModems in particular, see Chapter 9.)

Finally, Finally, Before You Get Started

The bootable first CD or DVD is the last thing you need before proceeding to install Linux. Fedora comes on the bootable DVD that comes with this book. If your computer doesn't have a DVD-ROM drive, see Appendix B for various other options.



If you have to make changes in your BIOS, be sure that you save them before you exit! It's easy to forget and end up wondering why the machine isn't doing what you told it to do.

Now, before you go any farther, there's one more thing I want you to do: Go to the following Fedora sites and bookmark them so you can easily go and reference them if you need help:



✓ **The Fedora Project Wiki:** <http://fedoraproject.org/wiki/>

A *wiki* is a Web page (or collection of Web pages) maintained collectively, where anyone can make changes (or a group of designated people can, depending on how the wiki is set up). Wikis are becoming popular for generating community-based support sites because anyone with the desire to contribute can pitch in with a minimum of hassle.

✓ **The Fedora Documentation Project:** <http://fedoraproject.org/wiki/DocsProject>

✓ **The Fedora Forum:** www.fedoraforum.org

✓ **Fedora News:** www.fedoranews.org

✓ **The Unofficial Fedora FAQ:** www.fedorafaq.org

✓ **LinuxQuestions.org:** linuxquestions.org

If you're ready to install, proceed to Chapter 3.

Chapter 3

Installing Fedora

In This Chapter

- ▶ Considering some last-minute issues before installation
 - ▶ Installing Fedora as a personal desktop
 - ▶ Booting Fedora for the first time
-

Do, or do not. There is no “try.”

— Yoda, *The Empire Strikes Back*

No longer are arcane glyphs and complex sorcerer’s spells required to install Linux. The graphical installation is now quite easy to perform and will be familiar to users coming from another graphical operating system, such as Microsoft Windows. This chapter provides the details.

The installer is actually pretty smart. You may find that you don’t see exactly the same screens that I show in this Chapter. If you see something unfamiliar or don’t see one that I cover here, don’t panic. The installer is just adjusting what it offers based on the hardware in your system and what you chose to install.

Things to Consider Before You Begin Installation

You can install Fedora by

- ✓ Booting with the DVD-ROM included in this book and choosing Install or Upgrade an Existing System from the initial menu.
- ✓ Booting with the Fedora installation CD, which is CD number 1, if you downloaded the CD images yourself. Go to <http://fedoraproject.org/wiki/> and click the Torrents link, look for Fedora Core 7 Final, and the entry with a description similar to “Fedora Core 7 binary CD iso

images for i386” or the one for x86_64, if you have a 64-bit computer. Alternatively, you can order them from a third-party vendor such as

<http://shop.cheapbytes.com/>.

To begin the installation from the DVD-ROM or CD-ROM, you may first need to change your system to start, or *boot*, from a DVD-ROM or CD-ROM — many systems today are configured to do this already, so you may not need to make any changes. In Chapter 2, I cover how to configure a computer to boot from these devices.

I concentrate on Fedora 7 in this book for two main reasons:

- ✓ Fedora 7 is included on the DVD with this book.
- ✓ Covering the installation of every Linux distribution in existence would make this book into a set of encyclopedias.

Also, I assume that you want to install a desktop. There are plenty of books out there that focus on servers, so my goal is that *Linux For Dummies* focuses entirely on those who want to use Linux as their actual desktop machine — in addition to a few cool “desktop support” functions you’ll find throughout the book, such as in the Part of Tens. It is just not possible to cover both desktop and server functions to a satisfying depth in a book of this size.

Please note that if you’re installing Red Hat Enterprise Linux, some other version of Fedora, or a different distribution of Linux, your screens *will* look different from what is shown in this book. Each Linux distribution’s installation routine covers the same basic tasks, but the specific actions may be presented in a different order, or be customized to look different on-screen. For example, one distribution may present account creation before network configuration, and another may have those two topics reversed. However, most distributions go through the same basic choices, so reading this chapter could still be helpful for something other than Fedora.



With Fedora, up until you reach the About To Install (GUI install) screen, you can back out of the installation without changing anything on your system. None of the configuration options you make *before* you get to the About To Install screen are saved to your disks. After you continue beyond the About to Install screen, data is written to disk, and your system is changed.

Getting Started with the Installation Process

In this section, I follow the graphical installation. If you can’t use the graphical installer for some reason (if Linux doesn’t support your video card, for example), follow the text-based installation instead. The steps are the same,

it's just not as pretty, and you don't quite get all the options that are available in the point-and-click version.

Those who are using the DVD that came with this book should do the following:

1. **If you want to boot and install from the Fedora DVD-ROM that comes with this book, place the DVD in your DVD-ROM drive, reboot your system, select Install or Upgrade an Existing System, and press Enter.**

There are a number of boot options, shown in Figure 3-1, which are each described in the following list:



Figure 3-1:
The initial
Fedora
installation
screen.

- **Install or Upgrade in an Existing System:** The first (default) option in the graphical interface is for installing Fedora for the first time, or for upgrading an existing version of Fedora Linux. The graphical interface is designed to work with a mouse to select options. If you don't have a mouse, you can use the keyboard to navigate around the screens. In most places, the Tab key or the arrow keys advance you to the next option; the spacebar toggles options off and on; and the Enter key accepts the choices and moves to the next screen. In most screens, if you want to change a previous setting, a Back button is available to navigate to earlier selection screens. If you want more information about the installation process, click the Release Notes button at any time.

Dealing with damaged CDs or DVDs

If your CD or DVD appears to have a problem, what you do next depends on where you got the DVD or CDs. If the DVD came with this book, contact Wiley Customer Care at 800-762-2974 to get the disc replaced. Do not contact Red Hat's technical support to have the DVD replaced in this case.

On the other hand, if you burned your own CDs or DVD, you may be experiencing one of

two different problems. First, the burn may have become corrupted; to solve this problem, try burning the disc again at a slower speed. If the newly burned disc also fails the media check, the file(s) may have become corrupted during download. To solve this, download the images again.

- **Install or Upgrade an Existing System (Text Mode):** The second option enables you to install or upgrade Fedora by using a text-menu interface. All the options appear in text menus, and you use the arrow keys or Tab key to move to the selection you want. In some areas, you use the spacebar to turn options off and on.
- **Rescue Installed System:** This option is not for installing. Instead, you use it to boot into *rescue mode* if there's something wrong. See Chapter 4 for more on rescue mode.
- **Boot From Local Drive:** This is also not an installation option. Selecting this menu item lets you boot from a hard drive in case your boot menu breaks after installation.

If you don't choose any options and just press Enter, the Fedora graphical installation process starts. I will assume you did this, which causes many lines of information to scroll past as the installer launches.

If you downloaded the full version of Fedora and burned it onto CDs or a DVD yourself, the CD Found (media-check) screen appears. This screen allows you to check the integrity of the media that you're using to install Linux. If this is the case, I recommend taking the time to perform this test on all the CDs or the DVD you downloaded personally. It's best to know now whether one of them is damaged or incomplete.

Whether you check your media or not, the initial Fedora installer screen eventually appears. If you want, click the Release Notes button to open the file containing information about this version. It may be a bit geeky but you might find something interesting or useful.

2. When you're ready to leave the initial Fedora installer screen, click **Next**.

The initial screen disappears and is replaced by the language-selection screen, as shown in Figure 3-2. Linux supports many different languages.




Figure 3-2:
The
language-
selection
screen.

3. Use the mouse or the ↑ or ↓ keys to select your language and click Next.

The keyboard-configuration screen appears, as shown in Figure 3-3.




Figure 3-3:
The
keyboard-
config-
uration
screen.

Different languages arrange the keys differently on keyboards; you may want to choose the matching language for your keyboard. (The default is U.S. English.)

4. Choose your keyboard configuration and click Next.

The partitioning screen appears, as shown in Figure 3-4. You're given four initial options:

- **Remove all partitions on selected drives and create default layout:** Wipe everything off the hard drives you select in the check boxes below this option.
- **Remove Linux partitions on selected drives and create default layout:** If you already installed a version of Linux on the hard drives marked with the check boxes shown below this option, the installer will find the previous instance and install over it, wiping the old one out.
- **Use free space on selected drives and create default layout:** Find available free space and install there.
- **Create custom layout:** Manually define where to install. If you go this route, see the nearby sidebar “Creating a custom layout.”



Figure 3-4:
The
partitioning
screen.

As a hint for choosing the proper drives, the first hard drive in your computer is typically `hda`, the second is `hdb`, and the third is `hdc`. (With some kinds of hard drive technology, the first drive will be `hde`, the second `hdf`, and so on — just count from the lowest to the highest.) If you have SCSI drives, you see `sda`, `sdb`, and `sdc` instead.

Creating a custom layout

If you choose to define partitions manually, you'll be using Red Hat's Disk Druid. Disk Druid enables you to delete existing partitions and add new partitions. If you don't know how to use Disk Druid, I recommend that you *don't* proceed any farther. Click Back to return to the Automatic Partitioning screen, and then, instead of starting from scratch, choose to review the automatic setup that the Disk Druid suggests.

If you choose to Review and decide to change things, here are some terms you'll want to know: The *mount point* is where a *device* (maybe the first partition on your first IDE hard drive, which would be `/dev/hda1`) attaches within your *filesystem* (your directory hierarchy). The start

for your filesystem is the *root directory*, which is shown as `/`. You need to have a partition mounted as `/`, and you want it to be as large as possible if you're not making many partitions. You also want a partition mounted as `/boot`, but that one can be small, around 150MB. Finally, you want a partition that's about twice as big as the amount of memory in your computer. (There's no mount point shown for the swap partition, just a file-system type of `swap` that you have to choose. Otherwise, when it comes to Fedora, you'll be sticking with `ext3` as your file-system type because it's the default. (I get more into these kinds of things in Chapter 10.)

If you're both nervous *and* technical enough to understand hard-drive lingo, click the Review and modify partitioning layout check box first. If you do that, a layout screen appears. Those who choose to delete existing partitions will be asked if they really want to do this. Refer back to your notes to double-check before you actually make anything go away.



Be careful — you lose *all* data in existing partitions when you choose to remove partitions. They won't be immediately changed, however — not until the final installation screen.

You will also find an Advanced Storage Configuration button. Concern yourself with this button only if you know what iSCSI and DMRAID are and need to deal with them.

5. Choose the appropriate option and make sure the correct drives are selected; then click Next.

After you click Next, what appears depends on what the installer decides you need to see. You may find yourself at the GRUB boot loader screen. In this case, it is probably that you're either deleting Windows or are setting up a dual boot. If you are having Linux and Windows share a computer, make sure they both appear in the list of options. Be sure to have the Default check mark attached to the operating system you want to boot into by default.

Otherwise, the network-configuration screen appears, as shown in Figure 3-5.

Figure 3-5:
The
network-
config-
uration
screen.

The screenshot shows the Fedora network configuration window. At the top is the 'fedora' logo. Below it is the 'Network Devices' section, which contains a table with columns 'Active on Boot', 'Device', 'IPv4/Netmask', and 'IPv6/Prefix'. The first row shows 'eth0' as the device with 'DHCP' and 'Auto' settings. An 'Edit' button is to the right of the table. Below the table is a 'Hostname' section with the text 'Set the hostname:'. There are two radio buttons: 'automatically via DHCP' (which is selected) and 'manually'. The 'manually' option has a text box containing 'localhost.localdomain' and a hint '(e.g., host.domain.com)'. Below the hostname section is the 'Miscellaneous Settings' section, which includes text boxes for 'Gateway', 'Primary DNS', and 'Secondary DNS'. At the bottom of the window are three buttons: 'Release Notes', 'Back', and 'Next'.

If you're using a modem and don't have an Ethernet card in this system, you won't see this screen at all, so skip to Step 7. If you do see this screen, go ahead with Step 6.

- 6. Configure your network. If your ISP or network Administrator told you to use the Dynamic Host Configuration Protocol (DHCP) to let the computer automatically get its network settings, most of the job is done for you. Otherwise, you will need to set the information manually.**

If your networking is being set through DHCP, the only item you may want to change in this screen is the *hostname*, which gives a name to your computer on the network. You can choose to override this name and set your Linux system's hostname manually (instead of letting the DHCP server pick one for you). To do so, follow these steps:

- a. Select the Manually radio button rather than Automatically Via DHCP.**

The hostname text box becomes active.

- b. Change the value in the text box by typing in a hostname.**

If you are not using DHCP, but instead are using *static* networking — which means you feed in specific values — do the following:

- a. Click the Edit button next to the Network Devices box.**

The Edit Interface dialog box appears.

- b. In the Enable IPv4 support section, click the Manual Configuration radio button.**

The text boxes below the check boxes become available.

- c. Take the IP address that your network Administrator gave you and type it into the IP Address text box.

For example, if your IP address is 192.168.1.46, you type it into the IP Address text box.

- d. Type the *netmask* (network “mask,” or address) given to you by your network Administrator into the Prefix (Netmask) box.

If you were given information that you were told is for IPv6 addresses, fill in the boxes under Enable IPv6 Support as well. Otherwise, just ignore this section.

- e. If you want your computer to connect to the Internet at boot time, make sure the Activate On Boot check box is checked; otherwise, click it to remove the check.

- f. Click OK to close the dialog box.

- g. In the hostname text box, enter the hostname your Administrator asked you to use.

- h. In the main screen, enter the values in the Gateway and DNS fields that your network Administrator gave you.

If your network has more than one DNS server, you can enter up to three DNS server addresses — a DNS server helps your computer find other machines on the Internet.

- 7. After you’ve set your configuration, click Next.

The time-zone selection screen appears, as shown in Figure 3-6.



Figure 3-6:
The time-
zone
screen.

8. Click your city, or if you don't find it in the listing, choose another city in the same time zone that supports the same options (such as daylight-savings time).
9. Click Next to proceed to the next step in your installation.

The root password screen appears.

10. In the root password screen, type the root (Administrator) account password into the Root Password field and then type the same password in the Confirm field.

You don't see the password when you type it — just a dot for each character. (The dots prevent unauthorized individuals from seeing the password.) If you mistype something in one of the boxes, you're warned when you try to move on to the next step of the installation and have a chance to reenter the values. You will also be warned if you chose a password the installer considers weak — and therefore a danger to your system's security.



Don't forget your root password! You need it to do any administrative task on your Linux box. If you do, though, see Chapter 22 for how to reset it.

11. After you've entered your root password twice, click Next.

The application types screen appears, as shown in Figure 3-7.




Figure 3-7:
The
application
types
screen.

12. Add or remove types of packages to customize your system.

To add a type, there needs to be a check in its check box. To remove one, the box needs to be empty. You can add and remove checks by clicking in the box.

The groups available are

- **Office and Productivity:** I assume you are setting up a desktop system, so you want to make sure this one is selected.
- **Software Development:** If you are a programmer, you want to add this group. Otherwise, you can do without it.
- **Web server:** If you want to run a Web server on this system or just experiment with the Apache Web server, add this one. Otherwise, you don't need it.

Keep in mind that the last two options are beyond the scope of this book.

13. **At the bottom of the screen, select *Customize Now*, and then click *Next to proceed*.**

The package selection screen appears, as shown in Figure 3-8.



Figure 3-8:
The
package
selection
screen.

14. **Click *KDE (K Desktop Environment)* if you want to have both of the major graphical interfaces installed on your system.**

I focus on GNOME (see Chapters 5 and 8) in this book but KDE is also popular among many Linux users.

15. **Click *Applications*.**

The types of packages available in this category appear on the right. You can start anywhere you want among the options on the left, but I thought this section might be a good example.

16. **Scroll through the list on the right. Make sure the check boxes are checked for the groups you want, and aren't checked for the groups you don't want.**

You can click the name of the group to see a description of it below.

17. **If you want to look through the optional packages (programs) for this group, click the Optional Packages button while it is selected.**

This button is not available if the group isn't already marked for installation. For each package group, there are programs that are installed by default, and then programs that are considered optional. You have to select the optional packages yourself in order to install them, and you do so in the Packages dialog box that opens when you click Optional Packages for that group. Read descriptions and add/remove checks from check boxes until you are satisfied with your choices, and then click Close to close the dialog box.

18. **When you have finished choosing optional packages for each group, choose the next category on the left (for example, Development for Software Programming Tools) and then return here to Step 18.**

19. **When you have finished going through all the possibilities — you can change what is installed later — click Next.**

First, a dialog box appears and tells you that the installer is looking over your list of selected software. If it discovers that you've left out programs that the software you chose depends on, it offers you the chance to add those too. (Say yes!) When all that's done, the install screen appears.



If you want to stop your installation of Linux and/or changes to your hard drive, the install screen is the *last* place where you can stop without changing anything on your hard drive(s). To stop the installation, press Ctrl+Alt+Del and your system reboots. If you do that, be sure to pull out the DVD or first CD if you don't intend to restart your installation.

20. **When you're ready to commit to the installation, click Next.**

The Required Install Media dialog box may appear if you're installing from CDs, telling you exactly what CDs you need for the installation. If you're using CDs, you may not need all of them, depending on the software you requested.

After this, the package-installation screen appears. The system first prepares for the installation and then starts installing. As the system is installing, you see the name of each individual package being installed, and a progress bar for the total installation.

After the package installation has finished, you reach the final installation screen, which has text beginning with “Congratulations, the installation is complete.” The CD or DVD is automatically ejected (if not, eject it manually). After you’ve removed the installation media, click Reboot to restart your machine and proceed to the next section of this chapter, “Your First Boot.” If you want to shut off the computer instead, it is safe to do so at this time. The next time you start it, you will proceed to the process discussed in the next section, “Your First Boot.”

Your First Boot

The first time your Fedora system boots, you have your first chance to see your boot menu fly by. It’s the blue screen that offers you three seconds before proceeding to boot into your default choice. If you set Windows as your default, press a key (such as the spacebar) to enter the menu, select the Linux entry, and press Enter. If Linux is the default — or the only — operating system installed, just sit back and let the machine boot on its own. If, for any reason, your computer fails to boot at this point, see Chapter 4.

The first time your computer boots, you see the first boot Welcome screen. Do the following to complete your machine’s initial setup:

1. **Click Forward to proceed to the setup routine.**

The License Agreement screen appears.

2. **Read through this text, select Yes, I Agree To The License Agreement and then click Forward.**

The Firewall screen appears, as shown in Figure 3-9.

Firewalls protect your system from unauthorized access and discovery. When you’re connected to the Internet at high speed, your computer becomes one of hundreds of thousands. I wish that I could tell you that there *aren’t* tons of nasty people and programs trying to break into your computer, but I can’t; unfortunately, there’s a horde of them. A firewall is your first (but not only) protective barrier against cyber-hoodlums, viruses, worms, and a host of other items that might try to skulk in through doors that should have been kept closed; see the sidebar “Firewalls and you” for more information.

3. **Unless you are 100 percent sure that this computer is safe from both the Internet and other machines you may have on your home or office network, leave Enabled as your firewall setting.**

It’s worth keeping in mind that before viruses and worms, most business computer break-ins happened from *inside* rather than outside. If you are absolutely determined to have your firewall off, select Disabled. I wish you luck!



Firewalls and you

Anyone connected to the Internet, especially with *broadband* (see Chapter 9 for more on the different types of connections available; broadband tends to be DSL and cable), really should have a firewall of some sort. You can purchase a firewall *appliance* (a special little box that you attach between your Internet connection and your computer[s]) or have a Linux box doing

your firewalling for just itself — or even for your whole home or business network. Although this is quite an extensive topic that can easily take up an entire book of its own, I cover some basics in Chapters 9 and 19. In the meantime, protect at least this machine with a firewall of its own.



Figure 3-9:
The Fedora
Firewall
screen.

4. Determine which of the available services you want to let through your firewall.

I recommend leaving SSH enabled so you can access this computer remotely, using the techniques discussed in Chapter 18. Although I don't cover running your own Web server in this book, you can run one on this computer — if so, you'll also want to enable the WWW (HTTP) and Secure WWW (HTTPS) entries (no need to enable these if you just want to browse the Web). The same goes for FTP (which can be handy for moving files around) and Mail (SMTP). If you intend to run an FTP server or a mail server (SMTP), you'll need to enable these services, but most desktop users won't need Mail because they're using their ISP's e-mail.

5. **Next to each item you want to enable, click in the box so it has a check mark in it; click Forward when you are ready to proceed.**

For my desktop machines, I tend to check the SSH and FTP entries. For how to use an FTP client to get files off this machine later, see Chapter 8. When you click Forward, a Warning dialog box may appear.

6. **If you see the Warning dialog box, click Yes to set the firewall configuration to what you have selected.**

The SELinux screen appears.

7. **Select Permissive unless you are positive that you want to experiment with SELinux (otherwise known as Security-Enhanced Linux).**

SELinux is explained in great geekly terms at

```
http://fedora.redhat.com/docs/selinux-faq/
```

This extra layer of security is quite an advanced topic, far beyond the scope of this book. In fact, it's new enough that even seasoned Linux Administrators are still trying to figure it out.

8. **Click Forward to proceed to the next stage of your post-install configuration.**

After a few seconds, the Hardware Profile screen appears.

9. **Select Send Profile so the Fedora Project knows what hardware Fedora is being used on, and then click Forward.**

The data is sent anonymously and ensures that they know what you're running and keep it in mind during development.

10. **Type the name you want for your personal account into the Username text box, and then your name into the Full Name text box.**

For example, `zorro` in Username and `Zorro the Dog` in Full Name.

11. **Enter your account password both in the Password and Confirm Password text boxes.**

This password is the one for your personal account and has nothing to do with your root (Administrator) account. Feel free to use a different password than you did for root. In fact, it's a good idea to do so. If you accidentally don't type the password the same both times, you're warned in Step 16 and returned to this step. Also, if you choose a weak password, you will be warned that it's weak. It's a good idea to replace weak passwords so no one breaks into your system.

12. **Click the Forward button to proceed.**

That's it! You've just survived the second Linux gauntlet! Your computer now brings you to a graphical login prompt. See Chapter 5 for instructions on what to do from here. Enjoy exploring Linux!

Chapter 4

Booting and Stopping Fedora

In This Chapter

- ▶ Understanding what happens when you turn on your Linux machine
- ▶ Identifying and isolating boot problems
- ▶ Shutting down safely
- ▶ Taking Linux off of your machine

I like work; it fascinates me. I can sit back and look at it for hours.

— Jerome K. Jerome

If you came here from Chapter 3, you likely just survived the first gauntlet of the Linux world: installing the operating system. I hope that booting for the first time worked well. If it did, and you decide that most of this chapter isn't for you, at least skip to the section near the end of this chapter, "Don't Just Turn Off the Machine!" Otherwise, if you're interested in finding out about what your machine does as it boots (this chapter can get pretty geeky!), read on.

Your Linux installation failing to boot properly is not necessarily an emergency. Much of this chapter is designed to help you deal with any problems you might run into. Before you curse Linux, remember that installing an operating system is no small task — and because many technical variables are associated with such an installation, many computer manufacturers insist on performing the task at the factory. (When's the last time you installed Windows yourself?) Companies such as Hewlett-Packard, Dell, and Wal-Mart (believe it or not) have been offering Linux preinstalled, so if you're too discouraged at this point, or are reading ahead and are too nervous to give it a shot, you can either use the Fedora LiveCD discussed in Chapter 2 or purchase a preinstalled machine from a vendor or a local computer builder.

To unravel the mystery of a Linux system that won't boot requires a bit of understanding of how your computer starts up. At the very least, being able to tell a repair person where the problem seems to be can make you feel a little less intimidated — and can even save you some money because it speeds up their troubleshooting.

Giving Linux the Boot

Face it: As enjoyable as the experience of staring at a dormant computer is, the real fun starts when you turn it on. As with any electronic device, opening the electron floodgate is the first step to fun. A computer, however, can do much more stuff than your toaster oven. Rather than warm up a simple heating element, your computer has to check all those gizmos that you (or the manufacturer) plugged into your computer's motherboard. After the initial power-up, the computer performs some simple hardware tests (called the POST, or Power-Up Self-Test) to determine whether those various components are working properly.

Checking all your hardware is just the beginning. Between the time you turn the computer on and the moment the glowing phosphor on your monitor prompts you for a login name, the computer is building itself an empire. If you listen and watch carefully, your computer and monitor show you signs of the boot process through bleeps, buzzes, whirring motors, clicks, messages on the monitor, and blinking lights.

Although you have heard the cliché “Rome wasn't built in a day,” the boot process goes fairly quickly. This is pretty amazing, considering that the architecture of an operating system makes Rome's look like a stack of cardboard boxes; each time you power up your computer, it must build its whole operating system in memory. (Remember that an *operating system* is the core software that makes your computer work.) This process can be broken into four main steps, which I discuss in the following sections.

Step 1: Power-On Self-Test (POST) leads to BIOS

The POST process really has nothing to do with the operating system. Your computer performs this step whether you're running Linux or another operating system, such as Windows Vista.

Some symptoms of a failed POST include

- ✓ An unusual series of long and short beeps
- ✓ Nothing displayed on the monitor
- ✓ No apparent activity other than the whirring fan on the power supply
- ✓ A puff of smoke or the pungent smell of burning electrical components from your computer case
- ✓ An error message, displayed on the monitor, indicating a hardware failure

If you encounter any of these problems, you have hardware troubles that need to be resolved before you can proceed. If your computer was running properly before you began your Linux installation, your computer *should* be getting through the POST just fine — POST problems don't tend to be caused by installing a new operating system; they're far more fundamental to the computer itself.

For all but the last of these errors, it's time to question your nephew Mortimer, who was last seen lurking around your computer with a screwdriver. (Or take your computer into a computer repair shop!) If you see an error message indicating a hardware failure, you might have a shot at fixing the problem in the BIOS. As the POST does its thing and finishes up, it briefly displays (usually at the bottom of the screen) instructions on how to enter “setup” (which refers to the BIOS). Typically, these instructions mention pressing the DEL key or a function key such as F1. When you press this key (provided all goes well), you usually see a blue screen with black-and-white text: This is the BIOS. From here, if you're familiar with your hardware, you can try to figure out and fix the problem. However, many people would rather have a root canal with no anesthesia than mess with this stuff, so you may want to grab the nearest teenager or computer-repair-shop wizard and entreat her to take a look.

The good news is that if you can get to the BIOS at all, the problem *may* be easier to fix than one where the computer fails before you reach the BIOS.

Step 2: The BIOS passes the baton to the boot loader

After the BIOS gives the okey-dokey with a successful POST, the BIOS locates the first hard drive in your system and reads the first *sector* of that disk, which is often referred to as the MBR (Master Boot Record). On that first chunk of disk is a small program called a *boot loader*. The boot loader doesn't know much about anything, except how to start loading your operating system.

Two boot-loader programs understand how to load a Linux operating system: LILO (LIinux LOader) and GRUB (GRand Unified Boot loader). LILO has been a tried-and-true boot loader for as long as Linux has been a gleam in a geek's eyeball; GRUB is newer — and much more sophisticated. Many distributions allow you to choose your boot loader. Typically, most people just stick with the defaults, which in the case of Fedora is GRUB.

Sometimes “bad things” can happen to your boot loader. One common symptom of boot loader problems is a message saying that the operating system could not be found. Another hint that the problem is your boot loader is if you never see your *boot menu*, which is a blue screen that comes up during the boot process that gives you three seconds to press a key and see a menu letting you choose what operating system and version to boot.

If you suspect a boot-loader problem, see the section “Entering Rescue Mode,” later in this chapter. This section helps you to fix the problem. If that doesn’t work, then depending on the errors, your hard drive may be damaged, or another piece of your hardware related to the hard drive may be broken. You may have to drag in a teenager or a computer shop if you’re not particularly computer-savvy. Or, try using the drive for something else like Windows and see if it works for that.



If you just installed Linux, try just reinstalling it if you find it won’t boot properly. Oddly, that can fix the problem.

Step 3: The boot loader (LILO or GRUB) loads the system kernel into memory

Every computer (and electronic device) requires an operating system to run. The technical term for this operating system, or at least the core of it, is the *kernel*. If the kernel is your maestro, orchestrating all the components of your computer and delegating resources in a logical and cooperative manner, the boot loader is the red carpet on which the maestro arrives.

The boot loader passes control over to the kernel when you see text like this:

```
Uncompressing Linux... Ok, booting the kernel.  
audit(1092566112,337:0): initialized
```

Of course, that text might fly past. That’s fine — just keep in mind that if you see this text and then soon see something like:

- ✓ The system freezes
- ✓ A few dots appear across the top of the monitor, and then the system freezes
- ✓ A few messages appear on the screen, and the final message reads
kernel panic
- ✓ The system reboots automatically

you may have a problem with the kernel. If the boot proceeds past that text, the kernel scrolls more text down your screen as it finds all of the system's hardware and loads what it needs to handle it. Anywhere in this stage, if you run into the list of problems given a moment ago, is again the realm of kernel failure (you might even see the message `kernel panic`).

Recovering from a kernel-loading failure or panic (aside from the dread operating system not found situation described in the previous section) is one of the more challenging issues in Linux. This task is, unfortunately, not for the faint of heart. On the other hand, unless you just installed the system, it's likely that the failed kernel is one that you got through an update. In that case, you're in luck. Write down any error messages you see. Then reboot, press a key at the blue screen to open the boot menu, and choose the next Linux option down from the top (default) one. That is the next latest kernel. Press enter and you should boot just fine (if not, reboot and try the third one down, and so on until you can boot). After you're up and running again, you have two options. You can either go to www.google.com/linux and search on the specific error messages you found to try to figure out what's wrong, or wait until the next kernel update comes out and hope that it won't fail to boot.

However, if it's your original installation that doesn't start properly, you don't have a fallback kernel option. If this is your problem, many people actually just reinstall and hope that this action fixes everything. Often, it does. Otherwise, try doing a search as mentioned earlier, or see Chapter 2 for places to look for help among the Linux community.

Step 4: Control is handed over to `init`

When the kernel finishes loading, it passes the start-up process to a program named `init`. The `init` program is responsible for starting all services and programs. You can see these programs launching as they scroll up the screen with `[OK]` or `[FAILED]` on the right side of the monitor. If you see these lines, you know that your kernel has finished loading successfully.

The main problem you may encounter with `init` involves *services* (programs that run in the background). One or more of these can fail to start properly, as indicated by the `[FAILED]` status shown during the boot process. Services usually fail because of misconfigurations or unsupported hardware drivers. Although it's not the end of the world if a service won't start, sometimes you end up having to wait a number of minutes before a really stubborn one finally admits to the system that it can't start. Just be patient and make a note of which service is giving you trouble. Later, you can turn to Chapter 18 and find how to shut that service off so it won't be a problem again.

Entering Rescue Mode

Heavy-duty system repair tends to happen in *rescue mode* — a special boot option. The benefit of this option is that you can perform necessary system surgery even if your installed system is hopelessly broken. The drawback is that this is a purely command line interface — you really have to know your stuff to find your way around.

Appendix A contains a list of Linux commands to try to help you out. Chapter 22 addresses, among other things, how to fix your boot loaders from inside rescue mode (if your problem is with the boot loader), and how to track down the error messages your kernel may have left behind. You may also be able to find helpful tips by reviewing your distribution's documentation and online help forums.

To enter rescue mode in Fedora (it's text based, prepare for typing!), place the DVD or your first CD into your DVD-ROM or CD-ROM drive and boot the machine. Then follow these steps:

1. **When the disk first loads, select Rescue Installed System from the menu and press Enter.**

This action begins booting the system into maintenance mode.

2. **Select your language and press Enter.**
3. **Select your keyboard type and press Enter.**

The rescue system does its thing for a while, perhaps a minute or two on a slow system.

4. **When asked whether you want to start the network interfaces, answer No unless you know you need to download something.**
5. **At the Rescue screen, select one of the three options offered, and then press Enter.**

Your three options are

- **Continue:** The rescue interface tracks down your installed Fedora Core system for you.

If you select this option (or the next) and it fails, you may need to reboot and restart the rescue system. After restarting, choose Skip.

- **Read-Only:** The same as Continue, but you can't make any changes to your hard drive installation.
- **Skip:** If you choose this one, it's like saying *Don't bother trying to locate the filesystem, just give me a prompt!*

I assume that you chose Continue. If so, a window appears giving you some information about what has happened. You will find the instructions



it give you echoed in the next Step, so don't worry about pressing Enter to close the window. You won't lose the information.

When you press Enter, a *shell prompt* (the rescue command line interface) appears, and you now have access to the rescue interface.

6. If the rescue process was able to load your Fedora installation, type `chroot/mnt/sysimage`.

Doing so enables you to use your system without having to type `/mnt/sysimage` in front of everything. As mentioned earlier, see Chapter 22 and Appendix A for suggestions on what to do from here.

7. When you're finished, keep typing `exit` until the machine reboots, and then eject the CD or DVD as the reboot happens.

Don't Just Turn Off the Machine!

Even when you're not tapping the keyboard or clicking buttons, Linux runs along in the background, doing lots of housekeeping chores. Some of these chores may involve swapping *cached data* to and from the fixed disk — a geeky way of saying that it's actually jotting down things onto your hard drive that before it had just been making sure to remember in RAM. When you shut off the power out of the blue, anything the computer was remembering but hadn't written is lost (kind of like when you or I fall asleep before making a to-do list for the next day).

If you're used to Windows 98, you need to recondition yourself: Don't just shut off the power when you're finished. For one thing, many Linux users leave their computers on when they're not using them — they just log out of their accounts so no one can mess with their stuff — and (oh yeah) shut off the monitor because monitors draw a lot of power.

When you do decide to turn off your Linux machine, for whatever reason, you must shut down the computer in an orderly manner. You can use one of these methods to shut down Linux properly:

- ✓ If you're in the GUI, log out of your account by choosing System ⇨ Shut Down — and then click the dialog box option that says (strangely enough) Shut Down.
- ✓ If you have a command prompt open, type the **halt** command at the shell prompt; when you press Enter, Linux shuts itself down and tells you when it's all right to turn off the machine.
- ✓ If you have a command prompt open, type the **reboot** command; when you press Enter, Linux goes through the motions of shutting itself down, and then immediately reboots the machine.

If you do accidentally cut the power to your Linux box, take heart; all is not lost. More often than not, you can reboot your computer and pick up where you left off. However, you may have to pay for your error by waiting during a quick file-system check while the machine makes sure that nothing was damaged. This process is similar to the one Windows uses when somebody powers off incorrectly: The operating system may need to scan its hard drives to make sure that everything is okay.



If you're plagued by brief power outages or spikes that cause your computer to reboot, look into getting a UPS (Uninterruptible Power Supply) for your computer. These heavy, rechargeable batteries are designed not only to protect your computer from damage (read the box to see what features a particular UPS offers) but also to provide an extra five minutes or more of power so you have a chance to close your files properly and shut down your machine when the power goes out. Better yet, for those really brief spikes, the only inconvenience you'll have to deal with if you have a UPS in place is listening to it complain with beeps. The spike won't affect your computer at all!

Removing Linux from Your System

Although I'd hate to see you exit the Linux universe, I'd also hate to see you get trapped in Linux if you don't want to use it! How to most easily get rid of Linux depends on what you want to do:

- ✓ If you want to replace Linux with Windows, just install Windows on the machine. It will overwrite all of Linux, including the boot menu.
- ✓ If you want to remove Linux from a dual-boot setup, boot into Windows, access a command prompt, and type **fdisk/mbr**. This action removes the Linux boot loader. It doesn't erase Linux, but you can then format the Linux drive(s) or partition(s) for Windows.



Windows XP and Xandros users have a special command they can use to remove this distribution. Boot your machine by placing your Windows XP CD-ROM into the Recovery Console. Change to the Windows directory, type **fixmbr** at the command prompt, and (oh yeah) press Enter.

Chapter 5

Checking Out the Desktops

In This Chapter

- ▶ Utilizing the command line
 - ▶ Getting into GNOME basics
 - ▶ Customizing GNOME
-

A bus station is where a bus stops.

A train station is where a train stops.

On my desk, I have a workstation . . .

— Steven Wright

A lot of people like to characterize Linux as a DOS-like environment, where all you can do is operate in this antique-feeling world where you have to type a lot of cryptic stuff without any pretty pictures. However, the Linux desktop offers you quite a nice working environment, as you will see throughout this chapter. The cool thing is that most of it is configurable. Those who like to customize their systems can have way too much fun changing things around.

Making the Best of the Command Line

The Linux command line interface might evoke the old DOS-like days to some, but it provides a quick, efficient way of entering commands and executing actions. Even if you're mostly a "point-and-click" person, after you get the hang of using the command line, you discover that it's faster to perform some tasks there than with a mouse in the GUI environment. However, if you prefer to use a GUI interface as your working environment, you can easily open a *terminal window*, which is a command-prompt window, to perform

your command-driven tasks without having to completely leave the point-and-click environment (see Chapter 6).

There are some interesting things that you might find useful to know about using the command line interface. One difference between the Linux command line and other interface command lines, such as the Windows MS-DOS prompt, is that Linux commands are case-sensitive. Typing `LS` is completely different from typing `ls`, and, in fact, gives you an error because there is no `LS` command. The Linux command line also has an *autocompletion* feature. If you know the first few characters of a command or filename, type part of it and press Tab to complete the rest automatically. For example, if you're trying to use the `less` command to view the contents of the file `/home/bob/grocery_list`, you can type `less /home/bob/gro` and press Tab to try and complete the filename. However, if you also have the subdirectory (folder) `/home/bob/group_projects`, you hear a beep. You can press the Tab key again, like double-clicking a mouse, to see the output:

```
grocery_list    group_projects
```

The cool thing here is that the command prompt beneath these items still has the text `less gro` so you don't have to retype it! You can then see that typing a `c` makes it clear that you're referring to `grocery` and not to `group`, so you just add a `c` to make `less groc` and press Tab again to finish the autocompletion. This technique can save you a ton of typing, especially with really long filenames!

Most Linux distributions also keep a running *history* of commands most recently used. To use this list, press the up-arrow key on your keyboard. This action pulls up the last command you typed. As you continue to press the up-arrow key, you step through the most recently entered commands from the most to least recent. If you accidentally pass what you're looking for, use the down-arrow key to go back. When the command you want appears at the command prompt, press Enter to execute the command, or edit it and then press Enter.

In Chapter 6, I cover the command-line environment in more detail.

The GNOME Desktop

GNOME stands for the GNU Network Object Model Environment — not that this expansion tells you much! Suffice it to say that GNOME is a full point-and-click environment. Figure 5-1 shows you what you see after you log in to your Fedora system.

Figure 5-1:
The default
GNOME
desktop in
Fedora 7.



To find out more about GNOME, visit the main GNOME site at www.gnome.org.

Keep in mind that the programs you have depend on the type of installation you chose and what customizations you made; if what you have is different from what you see in descriptions or the figures, don't panic!

The GNOME desktop environment is divided into four parts:

- ✓ The menus
- ✓ The menu-and-icon panel at the top of the screen
- ✓ The desktop panel on the bottom of the screen
- ✓ The icons on your desktop

The menus

GNOME has three primary menus, all of them visible on the upper panel. From left to right, they are

- ✓ **Applications:** Programs available through the Graphical User Interface (GUI). (See Table 5-1.)
- ✓ **Places:** Shortcuts to special locations on your hard drive. (See Table 5-2. Note that you don't *have* to use the special directories created for you in

places such as Documents, Pictures, and so on.) Chapter 8 addresses this menu and its uses in far more detail.

- ✓ **System:** Personal and system settings, along with more general overall system commands. (See Table 5-3.)



Menu items that have an arrow to the right offer submenus, which you can open by holding your mouse on that menu choice. Often the submenus have their own submenus within, offering even more programs.

Table 5-1 Applications Menu Contents, Listed in the Order That They Appear

<i>Menu Choice</i>	<i>What You Find</i>
Accessories	Small, specific-function programs. Contains a calculator, character map, dictionary, screenshot program, text editor, and more
Games	A collection of games
Graphics	A variety of graphics programs, including The GIMP (see Chapter 14)
Internet	A few Internet tools you will find discussed in Chapter 10, such as Evolution (e-mail and calendar program), Gaim (instant messenger), and Firefox (Web browser)
Office	The OpenOffice.org suite of applications discussed in Chapter 12 and other useful productivity programs
Sound & Video	Programs, such as a CD player and sound recorder, for working with your computer's multimedia hardware (see Chapters 13 and 14)
System Tools	Tools for managing, monitoring, and updating your system, such as a file browser and a software updater
Add/Remove Software	Opens the application installation and removal utility

Table 5-2 Places Menu Contents, Listed in the Order They Appear

<i>Menu Choice</i>	<i>What You Find</i>
Home Folder	The contents of your user's home directory, as discussed in Chapter 7
Desktop	The contents of your user's Desktop directory, which contains files and folders that should appear on your desktop

<i>Menu Choice</i>	<i>What You Find</i>
Documents	The content of your user's Documents directory, in which you can choose to place files you're working on
Music	The content of your user's Music directory, in which you can save your music files
Pictures	The content of your user's Pictures directory, where you can save your images
Videos	The content of your user's Videos directory, where you can save videos
Download	Your user's Download directory, where your online downloads end up
Computer	The hard drives and temporary media available on your system
CD/DVD Creator	A special folder for pulling together files that you want to burn onto a CD or DVD
Network	The computers and hard drives available on your network
Connect To Server	A tool for connecting to many different types of servers
Search	A file-system search tool that looks inside files as well as at their names
Recent Documents	Documents you have opened lately

Table 5-3 System Menu Contents, Listed in the Order They Appear

<i>Menu Choice</i>	<i>What You Find</i>
Preferences	Your individual user settings
Administration	Systemwide settings
Help	The GNOME help tool
About GNOME	Information about GNOME
About Fedora	Information about Fedora Core
Lock Screen	The capability to set your machine so no one can use your GNOME login without entering your password
Log Out <i>user</i>	Enables you to leave your current login session
Shut Down	Hibernate (put into sleep/hibernation mode), Restart (reboot), or Shut Down (shut off) the machine

The Lock Screen tool

If you have your screensaver turned on and choose the Lock Screen option from the Main Menu, your screensaver appears or fades to black. Then if anyone moves the mouse or uses your keyboard, a dialog box appears with your login name in it and a password field. You can get back to work by entering your password. Until then, you're safe in knowing that no one else can mosey up to your computer and send off a joke e-mail to your boss while pretending to be you. Note that if you're logged in as the root user, the Lock Screen option doesn't work.

This feature is on by default in Fedora. To turn it off, do the following:

- 1. Choose System→Preferences→Look and Feel→Screensaver.**

The Screensaver Preferences dialog box appears, as shown in Figure 5-2.




Figure 5-2:
The GNOME
Screensaver
Preferences
dialog box.

- 2. Click the Lock Screen When Screensaver Is Active check box to deselect it.**

- 3. When you're finished making changes, click Close.**

The dialog box closes.

In this dialog box, you can also change if the screensaver turns on or not, how long it takes the screensaver to turn on, and what screensaver to use.

You can also click the Power Management button to go to more settings.

The panels

Along the top and bottom of your GNOME desktop lives a pair of panels. Because these bars are neatly divided into sections, take a look at what's in each section from left to right, starting with the top panel. On the far left side of the top panel are the three menus discussed in "The menus" section earlier in this chapter.

As you continue your journey to the right, you run into a group of icons (refer to Figure 5-1). You can reach all these items through the menus, but they're on the Panel to make them easy to find. In Fedora, these icons are

- ✓ **Firefox:** The Worldplanet-with-a-mouse-pointer icon opens the Firefox Web browser. (See Chapter 10 for more.)
- ✓ **Evolution:** Just to the right of the Web browser button; launches an e-mail and calendar program. (See Chapter 10 for more.)
- ✓ **OpenOffice.org Writer:** Opens the OpenOffice.org Writer word-processing program. (See Chapter 12 for more.)
- ✓ **OpenOffice.org Impress:** Opens the OpenOffice.org Impress presentation creation program. (See Chapter 12 for more.)
- ✓ **OpenOffice.org Calc:** Opens the OpenOffice.org spreadsheet program. (See Chapter 12 for more.)

Then you see a large blank space where you can add new icons. After the blank space in Fedora, here's what you find, from left to right — though your particular installation may have extra options displayed:

- ✓ **Update applet:** This applet may or may not be visible, depending on whether there are updates available or not. Here you can see how many updates are waiting for you to download them. Just connect to the Internet, click the applet icon, and choose Apply Updates to download and install them.
- ✓ **User switcher applet:** If you've created more than one user account (root doesn't count here), then clicking your name in the upper panel drops down a list of the other users on the system. You can use this applet to switch temporarily over to another user account. If you haven't created other accounts, then clicking your name just shows your name.
- ✓ **Date and Time applet:** Here you can see the day and time, depending on how this applet is configured.
- ✓ **Master Volume Control applet:** Click this to open the master volume control.

Now for the bottom panel. From left to right, you find on this panel:

- ✓ **Hide/Restore Desktop Applications:** This button lets you minimize all of your running programs immediately and then reopen them again with just one click.
- ✓ **Taskbar:** In this large space, you find entries for each program running on your desktop. You can change a program's status by using the boxes as indicated:
 - If a program is minimized, you can open its window by clicking its panel task box.
 - If a program is maximized but buried under another program, click its task box on the panel to bring it to the front.
 - If a program is maximized and on top, you can minimize it by clicking its panel task box.
- ✓ **Workplace Switcher:** Allows you to work in four different desktop environments during a single login session. Each desktop environment has the same menus, panels, and background, but you can run different programs in each of the environments. It's an easy way to remain organized while you're working in multiple programs. Try it out. It's like having four monitors in one!
- ✓ **Trash Can:** A shortcut to your desktop trash can.



Want to move a program from one workplace window to another? Click the icon in the program's upper-left corner and select one of the following options:

- ✓ **Always on Visible Workspace:** Makes the window show up on all four.
- ✓ **Only on this Workspace:** Keeps the window from appearing in the other Workspaces (on by default).
- ✓ **Move to Workspace Right:** Slides the window horizontally to the right, into the "next door" workspace.
- ✓ **Move to Another Workspace:** Gives you the option of specifying Workspace 1 (far left), 2 (second in from the left), 3 (third in from the left), or 4 (far right).

You can see which workplace an open window is in by looking at the Workplace Switcher and spotting the little windows that match how your desktop is laid out.

Changing your panels

You can customize your upper and lower panels individually through the Panel menu. You open this menu by finding free space on the panel you want to work with, and then right-clicking to pull up the Panel's context menu. (For a list of what this menu's items offer, see Table 5-4.)



If you have so many programs open that you're using the whole width of the lower panel, you may not have any free space to right-click in. Try being very precise with your mouse to find somewhere. If you can't, you may need to close a program or two to clear space.

Table 5-4	GNOME Panel Menu Contents
<i>Menu Choice</i>	<i>What You Find</i>
Add to Panel	The dialog box that lets you add applets, menus, and other objects to your main panel
Properties	The options for setting this panel's behavior
Delete This Panel	The capability to delete a secondary panel but not the main icon panel
New Panel	The options for creating new panels that sit on different parts of the screen
Help	The Help browser for GNOME
About Panels	A dialog box with some basic Panel information

The Add to Panel dialog box

The Add to Panel dialog box (see Figure 5-3) is accessible from the Panel menu (see the previous section). This dialog box contains a list of *applets* — small, specialized programs you can use to add particular functionality to a panel. To add one of these applets to your panel, open the dialog box, select the applet you want to use, and then click Add. The applet now appears on your panel. If you right-click the applet and choose Move, you can then slide the applet along your panel until you have it where you want it, and then click to release it.

Figure 5-3:
The GNOME
Add to
Panel dialog
box in
Fedora.



Configuring an applet

After you have an applet placed and running, you may be able to play with configuration options. Some of these options enable you to change what information is displayed. Others have a variety of look-and-feel settings.

To check for which configuration and other options are available for your applet, follow these steps:

- 1. Right-click the applet and examine the shortcut menu that appears.**

This shortcut menu is different from applet to applet. The bottom portion is always the same: Remove from Panel, Move, and Lock to Panel (prevent the icon from moving). Common entries for the top portion are Help and About. The rest of the items are either configuration options (see Step 2) or special applet features, such as the ability to copy the date from the Clock applet.

- 2. Choose Preferences from the shortcut menu.**

Not every applet has a Preferences dialog box. If this one does, the dialog box opens when you choose this option, and whichever configuration features this applet has are displayed.

- 3. Alter the selections in the Preferences dialog box to customize this applet's behavior.**

Now you get to have some fun. Make changes so that you can see what this applet can do — as you make your changes, they appear in the

applet on your panel. Each applet has its own set of features, so I can't do a general walk-through here. Just remember that you can always go back and change the settings or remove the applet from your panel later.

4. **Click Close to save your changes and close the dialog box.**

Ditching an applet

You have room for only so many applets. And, if you're like me, you probably don't want to have every bit of free space cluttered with icons. To remove an applet from the panel, simply right-click the applet you want to remove and choose Remove from Panel from the shortcut menu that appears. With nary a whimper, the applet vanishes from the panel.

Adding a program to the panel or the Desktop

If you have a program you end up using often, you can add it to your panel by following these steps:

1. **Choose Applications and browse to the program you want to add to the bar.**

Don't open the program. Just point to the menu item with your mouse pointer.

2. **Right-click the program and choose Add This Launcher to Panel to add it to the panel, or Add This Launcher to Desktop to add it to the desktop.**

An icon for this program appears on your panel or Desktop.

After you have your program on the panel, you can run the program just by clicking its icon. If you added it to your Desktop, double-click the icon.



If you're not happy with where a desktop icon is placed, click it and then drag it to a new location. For a panel entry, right-click it and choose Move. This option allows you to drag the icon to where you want it. Click to fix it into place.

You can also add a *drawer* (menu button) to your panel for an entire menu. To do so:

1. **Click Applications and browse to the submenu you want to add.**
2. **Enter that submenu.**

For example, if you want to add a button for the Graphics menu to your panel, open the Main Menu, move your mouse to the Graphics menu, and then move your mouse to the right into the contents of the Graphics menu.

3. **Right-click to open the context menu.**
4. **Choose Entire Menu, and then either Add This As Drawer to Panel or Add This As Menu to Panel.**

A drawer shows you its contents in icons only, whereas a menu looks more like the actual submenu from the Main Menu.

Removing a program from the desktop or the panel

To get rid of an icon you have on your desktop, right-click the icon and choose Move to Trash from the shortcut menu that appears.

If you want to remove one of the programs on the panel, just right-click the icon you want to remove and choose Remove from Panel from the shortcut menu that appears. The icon vanishes from the panel. That's it!

Playing with GNOME desktop icons

Your initial desktop icons form a vertical line along the top left of your screen (refer to Figure 5-1). In order, here's what you see from top to bottom:

- ✓ **Computer:** Opens the Nautilus file manager (see Chapter 8) with a list of your CD-ROM drive(s), hard drive(s), and more.
- ✓ **Home:** Opens the Nautilus browser (Chapter 8) with your home directory's contents displayed.
- ✓ **Trash:** A GNOME shortcut that opens the Nautilus file manager to the Trash folder, which contains files that you dragged into it because you weren't sure if you wanted to delete them.

To use the trash can, drag into it any files you want to delete. Later, if you're sure that you want to be rid of them, you can empty the trash in one of three ways:

- ✓ Empty the entire contents by right-clicking the trash can icon and choosing Empty Trash from the context menu. When asked for confirmation, click Empty.
- ✓ Open the trash can by double-clicking the icon. Then delete the entire contents of the trash can by choosing File→Empty Trash.
- ✓ Open the trash can by double-clicking the icon. To delete an individual item from the Trash, right-click it to pull up the Trash context menu, and choose Delete from Trash from the pop-up menu. When asked whether you're sure, click the Delete button to finish the job.



You can select more than one item by holding the Ctrl key to individually select them even if they're not next to each other; holding the Shift key to select a range of items; or left-clicking and then dragging to collect all the items that are in a box together.

- ✓ Another way to remove items from the trash can is to delete them manually from `~/ .Trash`.

Adding themes

Many themes are available on the Internet for GNOME. I will warn you now, though, that you may have a bit of frustration at first in finding a theme that you like that also is compatible with your installation. Keeping this in mind, muster up your patience first. Then to find and grab these themes for your own use, follow these steps:

- 1. Point your Web browser to `http://themes.freshmeat.net`.**

You're taken to the Themes Web site, where any number of GUI customization items is offered.



If you're heavy into playing with customization, you can also check out `http://art.gnome.org`. In particular, check out `http://live.gnome.org/GnomeArt/Tutorials`.

- 2. On Freshmeat, click GTK, which is geek-speak for part of the GNOME desktop environment.**

This action brings you to the appropriate section, to ensure that you choose themes that work with your system.

- 3. Select the GTK 2.X Themes area.**

Now you're in the section for the latest version of GNOME.

- 4. Browse and choose the theme you want to try.**

If you create an account and log in, you can use the Sort Order drop-down list box to change the order in which the items are displayed. When looking over a theme's information, keep your eye out for the requirements (sometimes called "dependencies") that go with the theme. Some themes require additional "engines" (software that runs behind the scenes), and you want to avoid them if you're uncomfortable with finding and adding software at this point. You're safe if you're looking for GTK themes.

- 5. After you have chosen your theme, click its name to go to the theme-specific page.**

There, you can find any comments someone has posted about the theme.

6. Scroll down if necessary and click the link under Tar/GZ or Tar/BZ2, and then Save to Disk when asked what to do with the file, and finally OK to download the file.

7. Choose System→Preferences→Look and Feel→Theme.

This action opens the GNOME Theme Preferences dialog box, as shown in Figure 5-4.

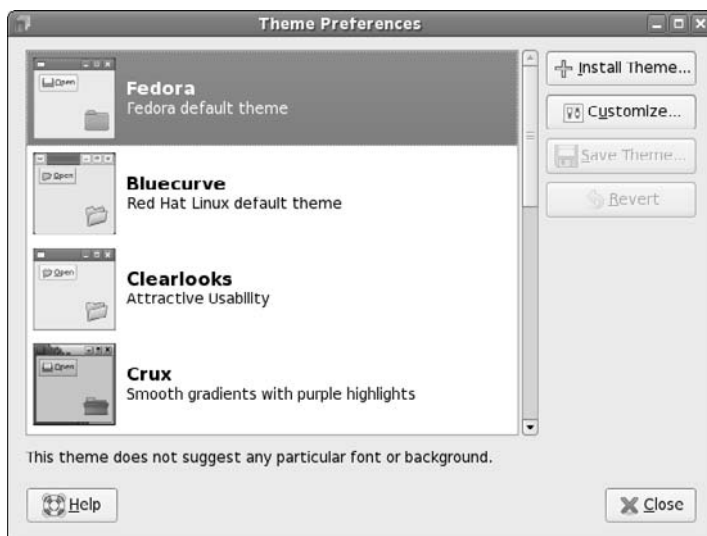


Figure 5-4:
The GNOME
Theme
Preferences
dialog box.

8. Double-click your user's Home directory icon on your desktop.

The File Manager opens.

9. Browse to where you stored the theme.

If you haven't changed the Firefox defaults, it will be in your Desktop directory.

10. Make sure that the theme file is a `.tar.gz` file.

If not, you need to see Chapter 16 on how to convert this file to this format.

11. Select the theme file and drag it into the left (Installer) pane of the Theme Preferences dialog box.

The system attempts to install the theme. If it succeeds, a confirmation dialog box appears.

12. Choose Apply New Theme if you want to change your theme right now. Otherwise, choose Keep Current Theme.

You may or may not actually see this question. Don't panic if you don't: There are different kinds of themes and some don't trigger the question. If you do see the question and you chose Apply New Theme, your desktop will change to match that theme. You can change to another theme if you want by clicking the theme in the list.

If you are told the theme is invalid, and the theme is in the specified format, you grabbed a theme that wasn't meant for your setup. If you look on the theme's page, it probably has something in its Environment or Dependencies section that tried to tell you that, but it's written for geeks, so don't beat yourself up. Just go back and try another theme!

Depending on the theme, your new theme may be added to the themes list. If you don't see it there, click Theme Details. In the Theme Details dialog box, you see three different types of themes: Controls for themes that affect how the controls on your desktop look, Window Borders for the themes that change how your window borders look, and Icons for themes that change how your icons look. Your new theme may be listed in one or more of these locations.

13. If you haven't already, select your new theme in the listing.

Whether you choose the theme in the main list or in the Theme Details dialog box, your desktop changes to match the theme. If the theme is not in the main list, you will see that Custom Theme is at the top of the choices. Click Save Theme and assign a new name to your desktop's new look.

14. Click Close when you're happy with your desktop's look, and the Theme Preferences dialog box closes.

Chapter 6

Working without the GUI

In This Chapter

- ▶ Understanding the user shell environment
 - ▶ Customizing the `bash` environment
 - ▶ Working with commands
 - ▶ Utilizing wildcard expansion
 - ▶ Using long commands
 - ▶ Tinkering with variables
 - ▶ Using redirection and pipes
-

Whom computers would destroy, they must first drive mad.

— Anonymous

Many computing old-timers speak fondly of the command line. Others who developed their skills by pointing and clicking refer to the command line as some antiquated tool used by crusty old-timers. The truth is that most skilled computing professionals recognize the merits of both the point-and-click graphical user interface (GUI), and the “lots of typing” command line interface (CLI). You must understand that the command line provides a powerful force for operating your computer. If you ever watch over the shoulder of a skilled Linux geek, you notice that, after logging in, the geek doesn’t take long to start tapping out seemingly cryptic instructions by hand.

In this chapter, I explore the Linux program that provides the CLI, which is called the `bash` shell. Although many shells are available for Linux, `bash` is the most common, and for good reason. Basically, the creators of `bash` rolled many good features of other shells into one terrific package.

Each shell has its own way of handling commands and its own additional set of tools. I start by explaining what a shell really is, and when you understand

that, you're ready to get down and dirty with `bash`. I cover specifically what you can do with some of the best features of the `bash` shell. Then I continue with working at the command prompt and get into `bash` shell interior decorating.



Shells come equipped to perform certain functions. Most of these features have evolved over time to assist the command-line jockey with myriad tasks. Although I only scratch the surface here, you're encouraged to read the `man` page for `bash` because it's likely one of the more complete and readable `man` pages in existence. You can read all about how to use `man` pages (the online Help system in Linux) in the "Help!" section, later in this chapter.

Playing the Shell Game

You need a way to tell the computer what you want it to do. In Linux, one of the ways to communicate with the computer is through something called the shell. A *shell* isn't a graphical thing; it's the sum total of the commands and syntax you have available to you to do your work.

The shell environment is rather dull and boring by graphical desktop standards. When you start the shell, all you see is a short prompt, such as a \$, followed by a blinking cursor awaiting your keyboard entry. (Later in this section, I show you a couple of methods for accessing the shell.)

The default shell used in Linux is the `bash` shell. This work environment is based on the original UNIX shell, which is called the Bourne shell and is also referred to as `sh`. The term `bash` stands for the *Bourne again shell*. The `bash` shell comes with most Linux distributions. You can start a `bash` session by selecting Applications→System Tools→Terminal. If you use your shell prompt quite a bit like I do, it's handy to use what you discovered in Chapter 5 to add the terminal window to your icons along the desktop's top panel.

Often, your shell prompt includes helpful information. For example, if you're logged in as `dee` on the machine `catherine` in Fedora 7, your prompt looks like this:

```
[dee@catherine ~]$
```

Before surveying a few of the shell capabilities, I need to tell you about another method for starting a shell session. You may not use it much but you might find it helpful at some point. First of all, notice that your shell prompt is merely inside a window that is part of your GUI desktop. Suppose that you want your whole screen to be text only.

To switch to a text environment, press `Ctrl+Alt+F2`. Don't be alarmed when your familiar graphical desktop disappears. It's still running in the background, and you can get back to where you left off in a moment. But first, a few words about the boring text screen you're looking at now (I hope).

You're looking at a virtual terminal, one of several available with your default installation. You probably see something like this:

```
catherine login:
```

Go ahead and type your username and password, which you're prompted for. You see a message indicating your last login date followed by the `bash` prompt:

```
[dee@catherine dee]$
```

Notice the similarity between this prompt and the open window you left behind in the GUI desktop. Both prompts are an indication that you have a `bash` session open. Note that, although it's accurate to say they're both the results of using the `bash` shell, they're distinct and separate *instances* of the same program. In other words, the environment you're working with here is exclusive of the `bash` environment you still have open in the GUI terminal window.

Are you wondering where your GUI desktop has gone? Just to settle your nerves a bit, do some jumping around. The GUI desktop is located at virtual terminal (VT) number 7 by default. You now have VT-2 open. Position your piano-playing fingers and strike the chord `Ctrl+Alt+F7` (if this doesn't work, try `Ctrl+Alt+F8`). Within a second or two, your screen should flash and return you to your graphical desktop. Neat, huh? And guess what? The `bash` session you left open on VT-2 is still there; you never logged out. Go back again by pressing `Ctrl+Alt+F2`. Voilà! — right where you left it. Feel free to jump back and forth a few times and try some other VTs (F1 through F6). Whoopee! This virtual terminal stuff rocks.



If you use the Lock Screen tool discussed in Chapter 5, only your GUI session is locked. Anyone could use `Ctrl+Alt` to move to a virtual terminal. So, make sure you log out of your virtual terminals before walking away from your computer.

Okay, when you have grown weary and bored with this little trick, exit (literally, type `exit`) to log out from each VT you may have opened and return to the graphical desktop and your `bash` prompt. Then you can explore what all the fuss is about with this *shell* doohickey.

Understanding bash Command Syntax and Structure

Many people happily skip through their Linux use without understanding the fundamentals of commands in the `bash` shell. Note that this approach makes you lose out on some cool capabilities available in `bash`. The more you know about how this shell's "language" works, the more interesting things you can do with it.

The basics of using `bash` at the command prompt often involve typing a command and any of its flags and values. For example, you enter the `ls -la ~` command to see a long-format listing of all files in your home directory, including those that start with a dot (`.`), which are *hidden files*. That other mysterious squiggle character is technically called a tilde. The *tilde* is a `bash` shortcut character that points to a user's home directory. For this example, I merely list the contents of my home directory.

You can break a command into three distinct components:

- ✓ The command name
- ✓ The options or flags
- ✓ The arguments

Consider this example.

Start with a simple command. The `du` command lists the contents of the directory you're now in, and its subdirectories, and how much hard drive space each item takes up, with a total at the end. Try typing just the `du` command by itself:

```
du
```

That's neat, but it probably raises more questions than it answers. The output gives you a long listing of data, but of what? Are those numbers in bytes, kilobytes, or messages from outer space? To clarify, try adding a simple option to your command:

```
du -h
```

You're still issuing the same command, but now you're providing additional direction on what you want displayed. The `-h` option tells `du` to show you the information in terms that humans can read more easily. Now *Ms*, *Ks*, and *Gs*

appear next to the numbers so that you can see how big these numbers actually are. But, wait — there's more. What if you just want to know the total amount of disk space this directory and its subdirectories are taking up? That calls for the `-s` flag:

```
du -s
```

What if you want the total for a different directory? Or just one of your subdirectories? In my case, I keep a Music subdirectory for the items I have copied from my CDs into Ogg Vorbis (see Chapter 13). I can type the following command to see how much hard drive space that directory takes up in a human-readable way instead of having to count zeroes:

```
du -sh ~/Music
```

In this example, `du` is the command name, `-sh` indicates the flags (options), and `~/Music` is an argument. The `-sh` flags can be accompanied by many more flags that provide various options applicable to the command.

Are you wondering where to find all the available options and arguments of a particular command? Most commands offer `man` pages, which are discussed in the “Help!” section, later in this chapter. Another good place to turn is the `--help` option, available with many commands. Note that `--help` displays a terse list of options, but it's nice and quick if you already know about an option but just can't remember exactly which one it is. Try it by entering the following command:

```
du --help
```

Cool, huh?

Starting Programs from the Shell

The most obvious, but perhaps not so apparent, use of the shell is to start other programs. Most utilities you use in Linux are separate and distinct executable programs. Users need a method to start these programs. In the GUI, you can associate an icon with a particular program, and the graphical environment contains the intelligence to start the program. Note that programs often require information drawn from environment variables, which are a part of the shell environment. (I discuss environment variables in more detail in the section “Working with Variables,” later in this chapter.) For this reason, the GUI often calls the intended program via the `bash` shell. So you see, even the GUI finds the shell a necessity — although the GUI does its best to hide this detail from users.

For example, in the GUI after you have a terminal window open, type the following command at the prompt:

```
mahjongg
```

After a few seconds, the Mahjongg game is displayed. You can start any program at a command prompt that you can click from the GNOME menu if you know what the underlying program name is. Note that if you're in a virtual terminal (press Alt+F1) rather than the GUI, you may see an error message. Some programs require a graphical environment in which to run, which a character-based terminal obviously doesn't have.

Putting Wildcard Expansion to Good Use

Computing life would be tedious if you had to repeat the same command on multiple files. After all, aren't repetitive tasks what the computer was designed to do? *Wildcard expansion* refers to the ability of one command to be executed against many files. The asterisk (*) and the question mark (?) are two wildcard characters that are used to match any filename, or a portion of a filename. For example, you can use the following command to see a long directory listing that includes only files that end with a .doc filename extension:

```
ls -l *.doc
```

The files listed may include `resume.doc`, `cover_letter.doc`, and `to_editor.doc`, for example.

Working with Long Commands

As you become used to the command line, you should find some shortcuts to ease your typing chores. In this section, I show you some features of the `bash` shell designed to make your life on the command line as pleasant as possible. These features include command-line completion, editing, and using the history of previously entered commands.

Asking Linux to complete a command or filename for you

Considering that you do much more typing on the command line in Linux than you may normally do in a GUI environment, a feature that provides typing shortcuts wherever possible is great. Command completion is a function of the shell that completes filename and system commands.

The capability of the Linux file system to deal with practically unlimited sizes of filenames means that many filenames can become huge. Typing these long filenames can become cumbersome. Fortunately, with command completion, typing a command or a long filename is short work.

You may want to use command completion in two situations: to enter a command or to complete a filename.

Completing a command

Suppose that you want to type a command, but you can remember only that it begins with the letters `up` and is supposed to return the length of time that has passed since the system was rebooted. Type `up` at the command prompt and then press Tab:

```
[dee@catherine dee]$ up[TAB]
```

One of two things happens:

- ✓ If only one matching command is in the *search path* (directory locations for searching for programs; type `echo $PATH` to find out what yours is), your command line is completed with that command, and the system waits for you to press Enter to execute the command.
- ✓ If you hear a beep, it means that more than one command begins with `up`. Simply press Tab a second time, and all the possibilities are displayed. Locate the command on the list and continue typing it until the first letters are unique, at which point you can press the Tab key to complete the command. If nothing is displayed after pressing Tab twice, there is no match at all.

Completing a filename

Command-line completion isn't only for commands; if you're typing a filename on your command line, you only need to type the first few characters and then press Tab. The shell usually searches the current working directory for filenames that match what you have typed and subsequently completes the filename on the command line. This feature behaves the same way as the

command-completion feature in that, if more than one file contains the letters you type, you hear a beep and need to press Tab again to see a list of choices.

It takes a little getting used to, but after you have control of the Tab key and the shell command-line completion feature, you may wonder how you ever got along without it.

Accessing your command history

It's nice of the shell to remember what you have done, for better or worse. Having the shell keep track of the commands you enter makes it easy to return to those gawd-awfully long commands you pecked at a while ago — even days ago! Let me give you an example. Suppose that yesterday you managed to issue a command to find all the *core dump* files in your system (core dump files are massive files containing debugging data that only an expert programmer or your computer can understand) and delete them. The command looked something like this:

```
find / -name core -exec rm {} \;
```

To reexecute the command, all you need to do is fish it out of your shell history and rerun it. The simplest way (if you're repeating the exact same version of the command you used last time, which in this case would be the `find` command) is to type `!find` and press Enter. Doing so tells your system to look through your history and rerun the last instance of `find` in the list.

On the other hand, if you have run the `find` command more than once and want to make sure that you're reexecuting the right version, you need to read through your command history. You can do so line by line by pressing the up-arrow key repeatedly until you locate the command you want to reexecute. Then just press the Enter key to run the command again.



The `history` command lists your last 20 commands (by default) when you enter it at the prompt, in case you're curious about what they were. If you want to see a longer list of them, use the techniques covered in Chapter 11 to view the contents of the file `~/.bash_history`.

Working with Variables

Variables in the `bash` shell are words or strings of text that computers use to represent a piece of data. An example of using a variable is setting the variable `fruit` to contain the text `apple`. A number of standard variables contain information about your account and environment settings.

Variables versus environment variables

The first thing I need to make clear is that the `bash` shell has two types of variables:

- ✓ **Variables:** A variable can be referenced in a program or shell session, but it's visible and available to only that session or program.
- ✓ **Environment variables:** An environment variable can also be referenced by the shell or program. However, it has the added behavior of having its value copied to any other program or shell that is created from its environment.



You can usually tell at a glance the difference between a variable and an environment variable in `bash`. The normal convention is to name local variables in all lowercase or in mixed-case characters. An environment variable, however, is usually always in all uppercase letters.

Checking out commonly used environment variables

The `bash` shell has many environment variables. You may be amazed at the range of items these variables store. The handy thing is that, if something is stored in a variable, you can change it to suit your needs! In Table 6-1, I list the environment variables you're most likely to want to work with.

Table 6-1 Commonly Used <code>bash</code> Environment Variables		
<i>Environment Variable</i>	<i>Purpose</i>	<i>Value</i>
HISTSIZE	Determines the number of previously typed commands that are stored.	Number of commands
HOME	Sets the location of your home directory.	The path to your home directory

(continued)

Table 6-1 (continued)

<i>Environment Variable</i>	<i>Purpose</i>	<i>Value</i>
MAILCHECK	Sets how often the <code>bash</code> shell checks for new mail in your mailbox. If mail has arrived, you see a message similar to <code>You have new mail</code> the next time you do something at the command prompt.	Number of seconds to wait between checks
PATH	Sets the directories that <code>bash</code> looks in, and the order to look in them to find a program name you type at the command prompt.	Colon-separated directories
PS1	Sets your command prompt.	Command and formatting characters used to form the prompt

Most environment variables are established for you by the system Administrator or perhaps by the shell itself. These variables are mostly read by programs to gather information, and you don't need to change their values. However, you may want to alter the value of some environment variables. For example, in Table 6-1, the first entry, `HISTSIZE`, determines the number of lines of command-line history that are kept on file. You may have read the discussion, earlier in this chapter, of reexecuting a command from yesterday. (If not, refer to the section "Accessing your command history.") By setting a higher number for `HISTSIZE`, you can save an even longer list of previously executed commands.

Storing and retrieving variables' values

To assign a value to a variable, you just use the variable name followed by an equals sign (=), followed by the value to store:

```
MyVariable=MyValue
```

To retrieve the value represented by that variable, you need to precede the variable name with a dollar sign (\$). Look at a variable, created by the shell, that determines what your prompt looks like. This variable is named PS1. First, you view the value being held by PS1:

```
echo $PS1
```

You likely see something like the following line:

```
[\u@\h \W]\$
```

Each of the characters preceded by a backslash represents a special instruction to the shell to return specific information when the shell prompt is referenced. See Table 6-2 for examples of special slash characters you can use in customizing your prompt.

Table 6-2		Pieces of the PS1 Puzzle
<i>Component</i>	<i>Result</i>	
\!	Prints the position of the command in your history list.	
\#	Prints the number of commands you have used during the current shell session.	
\\$	Prints a \$ for user accounts or a # for the superuser.	
\d	Prints the date in the following format: <i>day month date</i> .	
\h	Prints the name of the machine you're logged in to.	
\n	Moves down to the next line.	
\s	Prints <code>bash</code> for the <code>bash</code> shell.	
\t	Prints the time in 24-hour format.	
\u	Prints your username.	
\w	Prints the lowest current directory level.	
\W	Prints the entire current directory.	

Okay, on with the example; to change your shell prompt to something more amusing, enter the following line:

```
PS1='Hello \u, what can I do for you? => '
```

Note the single quotes. Immediately after pressing the Enter key, you see that your prompt has changed into something more inviting. Don't worry if you would rather have the original prompt: You can either reassign to the original prompt the value stored in PS1 or close the terminal window and open a new one, and you're back to familiar territory.

Are you wondering which other variables your system has in store for you? You can view all environment variables at one time by typing `env`. Note that you may not have any reason to access variables on the command line as a casual use of Linux. However, after you get more proficient, you may want to journey into the shell programming capabilities of `bash`, in which case variable storage is quite handy, just as it is in any computer programming language.

What's with those single quotes? You have to be careful of some details when changing environment variables. If you're just assigning something to a number, you could just use, for example, `HISTSIZE=250`. However, if you want to use something with spaces in it, you need to use quotes. Which kind of quotes you use depends on what else you want to do.

If you want to display *exactly* what you have specified, use single quotes to create a *literal text string*. For example, type the following line at a command prompt:

```
echo 'Hello, my name is $USER'
```

Kinda goofy, huh? Take a look at a different kind of string that the shell interprets differently: an *interpolated string*. An *interpolated* value is one in which the shell interprets special characters before processing the value. Rather than use single quotes, this time you use the same example with double quotes:

```
echo "Hello, my name is $USER"
```

Notice what the output is this time. Instead of displaying the exact text you provided, the shell replaces the variable name, designated with a dollar sign, with the actual value stored in that variable.



Why did I use single quotes in the PS1 example but double quotes with `echo`? The items with the backslashes (`\`) are *interpreted* one way or another. However, if you use double quotes with PS1, they're interpreted only once, so that item that lists what directory you're in changes only the first time. With a single quote, the variables are interpreted every time you do something. The double-quote example shows something like `Hello, my name is bob`. However, the single quote comes up with `Hello, my name is $USER`.

If you're going to play around with environment variables, I recommend that you start by using the methods I discuss in this section. After you have decided that you're comfortable with any changes you have made, you can make your changes permanent by opening the `~/ .bashrc` file and adding the same text there. The next time you log in, the changes go into effect. You can make changes for all your users' profiles in `/etc/profile` as well.



If you experiment heavily with these files, create a separate user account so that you can do whatever you want without messing up your own login. This advice especially goes for `/etc/profile`. You can damage everyone's logins with this one! To create a separate `/etc/profile`, you can make a backup by typing `cp /etc/profile /etc/profile.original`. Then edit `/etc/profile` knowing that you can always delete it with the `rm` command and use the `mv` command to rename `/etc/profile.original` to `/etc/profile`.

To create an environment variable from scratch, you will typically name it with all capital letters and then you have to “export” it. For example:

```
CUSTOMVAR="new variable"
export CUSTOMVAR
```

Don't be too discouraged if you don't understand all this variable stuff right now. As you become more proficient with Linux, you should explore *shell scripting*. Shell scripting is the art of creating computer programs with just the shell. Most Linux and UNIX Administrators speak shell script language like you and I speak our native tongues.

Using Redirection and Pipes

Redirection and pipes facilitate the flow of information. A *pipe* is exactly what it sounds like: It directs the output of one program to the input of another program. A pipeline may consist of several utilities plumbed together by pipes. At either end of this pipeline is, optionally, a redirection.

Almost all Linux utilities that require input and output have been plumbed with the following common interfaces: `stdin` (standard input), `stdout` (standard output), and `stderr` (standard error). By having a common method to feed input to a program or read data from the output of a program, you can glue utilities together into sophisticated solutions.

Redirecting command output

I discuss redirecting command output here because it's by far the most common form of information detouring. One example of *output redirection* involves telling a command to send its results to a file rather than to the screen, as you probably have been used to seeing. Start in some familiar territory by typing `ls -la ~` and then pressing Enter, to produce something like the following:

```
total 20
drwx----- 2 sue  users 4096 Oct 30 07:48 .
drwxr-xr-x 5 root root 4096 Oct 30 11:57 ..
-rw-r----- 1 sue  users  24 Oct 30 06:50 .bash_logout
-rw-r----- 1 sue  users 230 Oct 30 06:50 .bash_profile
-rw-r----- 1 sue  users 124 Oct 30 06:50 .bashrc
-rw-rw-r-- 1 sue  users   0 Jan  2 07:48 wishlist
```

Want to send this information to a file instead? You can use the `>` redirection operator to tell `bash` to send the data into a file rather than onto your screen. Enter the following command to send the information to a file named `listing`:

```
ls -la ~ > listing
```

Notice that nothing displays on the screen, as you normally would expect. That's because the shell has rerouted the output to a file named `listing`. To verify that the directory listing is there, enter the following command:

```
cat listing
```

The `cat` (and more) is explained Chapter 16.

Note that if you type `ls -la ~ > listing` again, the data is overwritten, meaning that the file's contents are wiped out and replaced with the new output. You can avoid this situation by using `>>` as your redirection operator, which tells `bash` to add the command's output to the end of the specified file. If you type `ls -la ~ >> listing` in the same directory after making no changes, the contents of `listing` are as follows:

```
total 20
drwx----- 2 sue  users 4096 Oct 30 07:48 .
drwxr-xr-x 5 root root 4096 Oct 30 11:57 ..
-rw-r----- 1 sue  users  24 Oct 30 06:50 .bash_logout
-rw-r----- 1 sue  users 230 Oct 30 06:50 .bash_profile
-rw-r----- 1 sue  users 124 Oct 30 06:50 .bashrc
-rw-rw-r-- 1 sue  users   0 Jan  2 07:48 wishlist
total 20
```

```
drwx----- 2 sue  users 4096 Oct 30 07:48 .
drwxr-xr-x 5 root root 4096 Oct 30 11:57 ..
-rw-r----- 1 sue  users  24 Oct 30 06:50 .bash_logout
-rw-r----- 1 sue  users 230 Oct 30 06:50 .bash_profile
-rw-r----- 1 sue  users 124 Oct 30 06:50 .bashrc
-rw-rw-r-- 1 sue  users   0 Jan  2 07:48 wishlist
```

Laying pipes

Another `bash` shell feature enables you to connect commands so that the output of one becomes the input for the next one. This feature is referred to as a *pipe*. Suppose that you want to look over the details of all files in the `/etc` directory in long-listing format. If you type `ls -la /etc` to do so, a massive listing appears, and much of the information scrolls right past you. Although you can back up a bit by pressing `Shift+PageUp`, you may not be able to see everything.

To see all the information, you can do one of two things:

- ✓ Send the data to a file with redirection by typing something like `ls -la /etc > ~/etclisting` and then review the contents of `~/etclisting` with your favorite editor.
- ✓ Pipe the output to the `more` command (see Chapter 16).

To pipe the output to `more`, type `ls -la directory_path | more`, where `directory_path` is the directory for which you want to list the contents. The `|` symbol (which on the keyboard looks more like two vertical bars stacked on top of each other rather than just one solid line) tells `bash` that you want to use a pipe.

“Help!”

The *man page* system is the electronic manual for Linux (*man* is short for *manual*), designed to provide users with a convenient reference to all the detailed command information. This information includes command-prompt options, file formats, and program function usage.

The syntax for opening a `man` page is `man <command name>`.

Don't know the command you're looking for or need basic information about using the `man` page system? Just type `man man` to get started. When you're finished reading the `man` page, press `Q` to exit.

Clearing the screen

The `clear` and `reset` commands are handy to know about when you're working in a shell. The `clear` command simply wipes the `bash` screen clean. Don't worry; it's not deleting any files or changing any settings — it's just tidying up so that you can start dumping new stuff to the screen again.

The `reset` command is a little more interesting. Suppose that you try listing a binary file to the screen with the `cat` command. After the computer finishes puking the result of executing the `cat` command on a binary file, you may get lucky and still be able to read your prompt. More likely, your prompt has been rendered into box characters of no special meaning, and typing on the keyboard gives you more of the same. To get back to normal, just type `reset` and press Enter. Note that it doesn't look like you're typing the word `reset`, but rest assured that the computer understands the series of characters and, after a couple seconds, should restore your shell environment to your native language.

Chapter 7

Getting to Know the Linux Filesystem

In This Chapter

- ▶ Discovering your way around the filesystem
 - ▶ Discerning partitions from directories
 - ▶ Adding removable media to the filesystem
 - ▶ Understanding filesystem permissions
 - ▶ Accessing shared files on Windows computers
-

I have an existential map. It has “You are here” written all over it.

— Steven Wright

One of the most frustrating things about mastering a new operating system can be figuring out where it keeps files. Instead of keeping all important system files in a single directory, such as the `C:\Windows` directory in Microsoft Windows, Linux follows the lead of its UNIX cousins and spreads things out a bit more. Although the Linux and Windows setups involve different methods, they are both logical, although it may not feel that way until you understand where to look.

After taking you on a tour of where to find things, this chapter shows you how to work in your filesystem on the command line. Reading this section isn't mandatory if typing in commands gives you heart palpitations, but it can be handy if you need to fix something later, and some people just like to know how to do these things.

Pieces of the Puzzle

It helps to understand the lingo before getting started. A lot of this will be familiar to you from other operating systems such as Microsoft Windows, but there are some differences that you'll need to get used to as well. To start, the word “filesystem” is actually used in more than one way. The general use (and what I am typically referring to throughout the book when I use this term) means “the files and directories (or folders) you have access to right now.”

The first major difference to be aware of is that Linux uses a forward slash (/) between directories, not the backslash (\) that Windows does. So, the file `yum.conf` in the directory `etc` is `etc/yum.conf`. Files and directories can have names up to 256 characters long, and these names can contain underscores (`_`), dashes (`-`), and dots (`.`) anywhere within, even more than one. So `my.big.file` or `my.big_file` or `my-big-file` are all valid filenames.

Upper and lowercase also matter quite a bit. They have to match exactly. The file `yum.conf` and `Yum.conf` are not the same as far as Linux is concerned. Finally, the same filesystem can span over multiple partitions, hard drives, and media such as CD-ROM drives. You just keep going down through subdirectories, not having to care if something is on disk A, B, or whatever. This issue will be addressed in the section “Partitions versus Directories” later in this chapter.

Touring the Linux Filesystem

In this section, I take a look at the contents of the typical Linux filesystem. Being at least a bit familiar with it helps you track things down later, and will help you to know where it's safe to mess with things and where you need to be very careful.

The root of the tree

Everything in the Linux filesystem is relative to the *root directory* — not to be confused with the system Administrator, who is the root user — which is referred to as `/`. This is the file-system base — a doorway into all your files. The root directory contains a mostly predictable set of subdirectories. Each distribution varies slightly, but certain standards exist to which they all conform. The standards keep the Linux world somewhat sane.



If you're interested in these standards, go to www.pathname.com/fhs and look at the latest version of the rules.

Table 7-1 lists what you might find in the base directories, meaning the items you find in /. (This list can vary depending on what you installed.) An asterisk (*) at the end of a description indicates that you shouldn't mess with this directory unless you have a *really good reason* because it contains files that are *very important* to the functioning of your system. Really, most of the base directories should be left alone. It's sections inside them that are safe to change.

Table 7-1	Standard / Contents in Linux
Directory	Contains
/bin	Essential commands that everyone needs to use at any time*
/boot	The information that boots the machine, including your kernel*
/dev	The device drivers for all the hardware that your system needs to interface with*
/etc	The configuration files for your system*
/home	The home directories for each of your users
/lib	The <i>libraries</i> , or code that many programs (and the kernel) use*
/media	A spot where you add temporary media, such as floppy disks and CD-ROMs; not all distributions have this directory
/mnt	A spot where you add extra filesystem components such as networked drives, items you aren't permanently adding to your filesystem but that aren't as temporary as CD-ROMs and floppies
/opt	The location that some people decide to use (and some programs want to use) for installing new software packages, such as word processors and office suites
/proc	Current settings for your kernel (operating system)*
/root	The superuser's (root user's) home directory
/sbin	The commands the system Administrator needs access to*

(continued)

Table 7-1 (continued)

Directory	Contains
/srv	Data for your system's <i>services</i> (the programs that run in the background)*
/sys	Kernel information about your hardware*
/tmp	The place where everyone and everything stores temporary files
/usr	A complex hierarchy of additional programs and files
/var	The data that changes frequently, such as log files and your mail

Some of these directories have some equally important subdirectories, which I cover in the upcoming sections.

The importance of being /etc

The /etc base directory contains primarily configuration files. It is best not to mess with these unless you know what you need to do. Some important directories here (which will vary depending on what you have installed) are

- ✓ The directories beginning with /etc/cron contain instructions for various programs that run at automated times.
- ✓ The /etc/cups directory contains configuration information for your printer.
- ✓ The directories beginning with /etc/rc all contain data about what starts and what doesn't at boot time and shutdown. When you see talk about *services* such as in Chapter 18, these are the directories that store when they run.
- ✓ The /etc/sysconfig directory has network configuration information.
- ✓ The directories beginning with /etc/yum have data about the software repositories (see Chapter 16) you use to add programs to your machine.
- ✓ The /etc/X11 directory contains configuration details for the X Window System (X), which runs your graphical user interface (GUI). See Chapters 5 and 8 for more on the GUI.
- ✓ The /etc/opt directory contains configuration files for the programs in the /opt directory, if you decide to use it.

Where temporary media lives

You may or may not have any subdirectories in `/media` or `/mnt` by default. Typically, however, you can look for the following:

- ✓ The `/media/floppy` directory is used for adding a floppy disk to your filesystem — rather than the word `floppy`; you might see `fd0` or `floppy0`.
- ✓ The `/media/cdrom` (or `/media/dvd`, `/media/cdrecorder`, and so on, depending on what type of hardware you have) directory is used for adding a CD-ROM, DVD-ROM, CD-Writer, and so on to your system.

In the “Managing Your Filesystem without a Net (Or Mouse)” section later in this chapter, I show you how to add these items on the command line. In Chapter 8, you find out how to do it in the handier point-and-click manner.

Where the computer is the /usr

The `/usr` directory is often referred to as its own miniature filesystem tree. This directory has lots of important or interesting subdirectories, as shown in Table 7-2. An asterisk (*) at the end of a description indicates that you need to leave that directory alone unless you have good reason to mess with it — *after* you gain lots of experience with Linux and know exactly what changes you need to make — so that you don’t accidentally alter something your system needs in order to function correctly. An interesting thing to remember about this section of the filesystem is that many advanced Linux users often use `/usr` to store programs that can be shared with other Linux machines.

Table 7-2 Standard /usr Subdirectories	
<i>Subdirectory</i>	<i>Contents</i>
<code>/usr/X11R6</code>	The files that manage the X Window System (the wire-frame underneath your GUI)*
<code>/usr/bin</code>	The commands that aren’t essential for users but are useful*
<code>/usr/games</code>	The games that you install on your system, except for those that you can choose to place in <code>/opt</code>
<code>/usr/include</code>	The files that the C programming language needs for the system and its programs*

(continued)

Table 7-2 (continued)

<i>Subdirectory</i>	<i>Contents</i>
/usr/lib	The shared code used by many of the programs in this /usr subhierarchy*
/usr/local	The programs and other items that you want to keep locally, even if you're sharing everything else in /usr
/usr/sbin	The commands that aren't essential for Administrators but are useful*
/usr/share	The information that you can use on any Linux machine, even if it's running incredibly different hardware from what this one is running*


Managing Your Filesystem without a Net (Or Mouse)

Chapter 8 covers how to use the point-and-click file managers. In this section, I take a look at how to use the command-line tools to find your way around the filesystem. It isn't vital for you to know this, but it could certainly prove useful someday!

Viewing information about files on the command line

As you may remember from Chapter 6, the `ls` command lets you view files and directories. Just by itself it shows you the names of the non-hidden files and directories that you ask for. For example, if you were in the `/etc/sysconfig` directory and typed `ls`, you might see what is shown in Figure 7-1.

Figure 7-1:
The multi-colored output of the `ls` command within `/etc/sysconfig`.



Links 101

A *link* is a fake file listing that points to another file, making a kind of shortcut. You use two kinds of links in Linux and UNIX: *soft links* and *hard links*. A soft link is like a Windows shortcut in that the link points back to the original file, and any edits you make to the link happen to the original file. Erase the original file, and the link remains, but it becomes unusable (and is called a broken link). The link is broken without the

original file. A hard link doesn't have a counterpart in the Windows world. A hard link isn't just a shortcut; it's another instance of the file itself. The data in this file is saved in only one place, but you can edit either the original or the link, and the edit is saved for both instances of the file. Erase the original, and the file still exists as long as the link is there. It's like two doors to the same room!

The colors displayed when you use this version of the command tell you something about the type of item you're looking at:

Text Color	Usually Means
Black	A regular file that doesn't match any of the special types that needs a color.
Red	An <i>archive</i> or <i>package</i> (see Chapter 16).
Pink	An image file (if it recognizes the file extension, such as .jpg, .tif, and .gif). Sometimes pink is also used for another type of special file, but it's heavily specialized and not something you need to worry about. If you don't see an image extension on the filename but you see a pink file, just ignore it.
Dark Blue	A directory.
Light Blue	A <i>soft link</i> (see the sidebar "Links 101" for what links are).
Green	A program (called an <i>executable</i>) in Linux.
Gray text with a red background	A <i>SetUID (SUID)</i> file (see the sidebar "Playing pretend with SUID and SGID" for what these terms mean).
Dark text with an orange background	A <i>SetGID (SGID)</i> file (see the sidebar "Playing pretend with SUID and SGID" for what these terms mean).
White text with a red background	A <i>broken link</i> (see the sidebar "Playing pretend with SUID and SGID" for what this term means).

Yellow text with a
black background

Black text with a
green background

A *device* file (basically, the file that is used to talk to a piece of hardware, like your monitor).

A directory where you can only delete or change the files that you own or have *write permissions* for (see the section “Understanding file listing information” later in this chapter for more).

The `ls` command shows everything in the requested location except for *hidden files*, which start with a dot. To see the hidden files too, type `ls -a`, and you get results as shown in Figure 7-2.

Figure 7-2:

The difference between a brand-new home directory with just `ls`, and then again with `ls -a` to see the hidden files as well.

```

[dee@localhost ~]$ ls
Desktop
[dee@localhost ~]$ ls -a
.      Desktop  .gnome2  .lessht  .redhat
..     .dmrc      .gnome2_private  .local   .themes
.bash_history  .egg cups  .gstreamer-0.10  .mcop    .Trash
.bash_logout  .esd_auth  .gtkrc-1.2-gnome2  .metacity  .wapi
.bash_profile  .gconf     .ICFauthority      .mozilla   .xsession-errors
.bashrc        .gconfd    .icons              .nautilus
.beagle        .gnome     .kde                 .qt
[dee@localhost ~]$

```

So far all we’ve done is look at files, and find out a little about them from the colors we see. The next section shows you how to discover more.

Playing pretend with SUID and SGID

Sometimes a program needs to be able to do things that the user running it doesn’t actually have the filesystem permissions to do. In that case, a file is allowed to pretend to be someone else in order to do what it has to do. A file that pretends to belong to another user when you

run it is said to have its SetUID (SUID, or Set User ID) permission turned on. Files that pretend to belong to a group that the user running them doesn’t belong to are said to be SetGID (SGID, or Set Group ID).

Understanding file listing information

To see more information than just names and vague types, you'll want a long format file listing. You get this listing with the command `ls -l` (if you want long format and hidden files, then you would type `ls -la`). Figure 7-3 shows a sample of what you might see.

Figure 7-3:
The beginning of the brand-new home directory with both hidden and non-hidden files displayed in long format.

```
dee@localhost:~$ ls -la
total 264
drwx----- 19 dee dee 4096 2007-04-04 23:07 .
drwxr-xr-x 3 root root 4096 2007-03-22 19:26 ..
-rw-r--r-- 1 dee dee 33 2007-02-12 07:18 .bash_logout
-rw-r--r-- 1 dee dee 176 2007-02-12 07:18 .bash_profile
-rw-r--r-- 1 dee dee 124 2007-02-12 07:18 .bashrc
drwx----- 6 dee dee 4096 2007-03-22 19:43 .beagle
drwxr-xr-x 2 dee dee 4096 2007-03-22 19:42 Desktop
-rw----- 1 dee dee 24 2007-03-22 19:42 .dmrc
drwxr-xr-x 2 dee dee 4096 2007-03-22 19:42 .eggccups
drwx----- 4 dee dee 4096 2007-03-22 19:44 .gconf
drwx----- 2 dee dee 4096 2007-04-04 23:35 .gconfd
drwxrwxr-x 3 dee dee 4096 2007-03-22 19:42 .gnome
drwx----- 6 dee dee 4096 2007-03-22 19:43 .gnome2
drwx----- 2 dee dee 4096 2007-03-22 19:42 .gnome2_private
drwxrwxr-x 2 dee dee 4096 2007-03-22 19:44 .gstreamer-0.10
-rw-r--r-- 1 dee dee 85 2007-03-22 19:42 .gtkrc-1.2-gnome2
-rw----- 1 dee dee 165 2007-03-22 19:42 .ICEauthority
drwxrwxr-x 2 dee dee 4096 2007-04-04 23:07 icons
drwxr-xr-x 3 dee dee 4096 2007-03-22 18:17 .kde
drwx----- 3 dee dee 4096 2007-03-22 19:42 .metacity
drwxr-xr-x 3 dee dee 4096 2007-03-22 19:43 .nautilus
drwxr-xr-x 2 dee dee 4096 2007-02-14 13:39 Pictures
```

You may find some parts of this format easier to understand, at a glance, than others. The first item in each listing (the part with the letters and dashes — for example, the `drwx-----` in the first line) is the *permission set* assigned to the item. Briefly, permissions define who can read the file, change it, or run it if it's a program. You can read more about permissions in “A permissions primer,” later in this chapter. The second item in the first line (in this case, 2) is the number of hard links to the item.

The third item (`dee`) is the file's *owner*, and the fourth (`dee`) is the *group* — depending on which version of Linux you're using, both these items may or may not be identical. You can find out more about both of these in “A permissions primer,” later in this chapter. The fifth item is the file's size in bytes. All directories show up as 4,096 bytes. Everything else has its own size. You can tell an empty file from the size of 0 bytes.

The sixth, seventh, and eighth entries are all related to the last time the file was changed: the month (`Dec`), the date (`3`), and the time in 24-hour format (`07:07`). Finally, the ninth item is the filename (for example, `bash history`, in the third row).

A permissions primer

If you find yourself scratching your head when looking at parts of that long format file listing, don't worry. The "Comprehending file types" section, later in this chapter, gives you a feeling for the first letter on each line, but nine more characters are attached to that item before you get to the next column. This group of nine is the set of *permissions* (also called a *permission set*) for the file or directory. Linux, UNIX, and even Mac OS X use permissions as a way of providing file and directory security by giving you the means to specify exactly who can look at your files, who can change them, and even who can run your programs. You need this capability when you have a bunch of different users on the same machine, networked to the world.

Checking out the triplets

Each permission set consists of three triplets. Each of the triplets has the same basic structure but controls a different aspect of who can use what. Consider the long format listing for `/home/dee` in the following code:

```
total 464
drwx----- 23 dee dee      4096 Dec 15 05:01 .
drwxr-xr-x  3 root root    4096 Dec  3 06:27 ..
-rw-----  1 dee dee         5 Dec  3 07:07 .bash_history
-rw-r--r--  1 dee dee        24 Jul 12 00:11 .bash_logout
-rw-r--r--  1 dee dee       176 Jul 12 00:11 .bash_profile
-rw-r--r--  1 dee dee       124 Jul 12 00:11 .bashrc
```

The first character in the permission set refers to the type of file. For a directory, the character is shown as a `d`, as you see here for the first two items in the preceding list; files are designated with a dash (`-`) instead. Each file or directory's permission set is a group of nine characters — that is, the nine characters that follow the first character (for a total of ten). But this group of nine is really three groups of three, as shown in Figure 7-4.

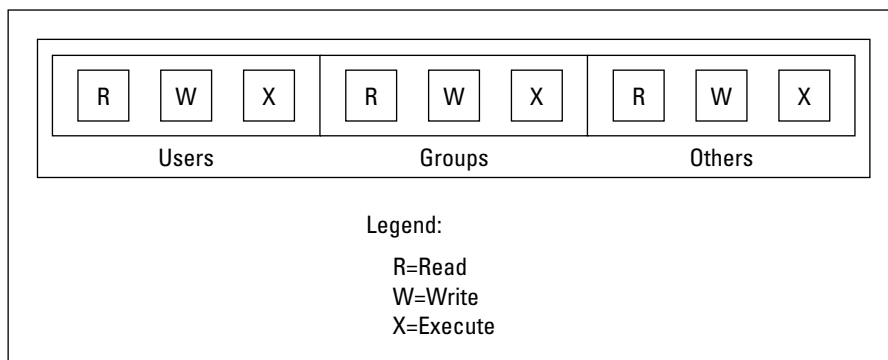


Figure 7-4:
Breakdown
of the nine
permission
characters.

The three triplets are read as follows:

- ✓ The first triplet consists of the second, third, and fourth characters in the long format file listing. This triplet sets the permissions for the *user*, or *owner*, of the file. (Owners are discussed in the “Beware of owners” section, later in this chapter.)
- ✓ The second triplet consists of the fifth, sixth, and seventh characters in the long format file listing. This triplet sets the permissions for the *group* that is assigned to the file. (Groups are discussed in the “Hanging out in groups” section, later in this chapter.)
- ✓ The third triplet consists of the eighth, ninth, and tenth characters in the long format file listing. This triplet sets the permissions for *other*, or everyone who isn’t the file’s owner or a member of the owning group.

Although each triplet is often different from the others, the internal structure of each one is made up in the same way. Focus specifically on how to read one triplet before looking at the set of them together. Each triplet includes three characters:

- ✓ The first character is either an *r* or a dash. The *r* stands for *read* permission. If *r* is set, the triplet allows the entity it stands for (user, group, or other) to view the directory or file’s contents.
- ✓ The second character is either a *w* or a dash. The *w* stands for *write* permission. If *w* is set, the triplet allows the entity it stands for to add, delete, or edit items in this directory or file.
- ✓ The third character is either an *x* or a dash. The *x* stands for *execute* permission. If *x* is set, the triplet allows the entity it stands for to access the files contained in this directory or to run the particular program in this file.



In all cases, if the dash sits in place of *r*, *w*, or *x*, the triplet doesn’t allow the entity the read, write, or execute permission.

The following sections describe owners and groups in more detail.

Beware of owners

You may have noticed by now that I talk a great deal about owners (users) and groups in Linux. Every file and directory has both of these components: a user from the `/etc/passwd` file that’s assigned as its owner and a group from `/etc/group` assigned as the group.

Although an everyday user probably doesn’t need to change file ownerships often, the root user does so regularly. If you add the file `comments`, for example, to `/home/tom` while you’re logged on as the *superuser* (another term for the Administrator, who is the person who owns the root account), root owns that file. The user `tom` can’t do anything with it unless you have set the last

triplet's permissions to allow the *other* folks (those who aren't the file's owner or in the specified group) to read and write to the file. But this method is a pretty sloppy way of doing things because the whole idea of permissions is to reduce access, not to give everyone access. Instead, remember to change the file's owner to the user `tom`. You do this with the `chown` (*change owner*) command. For example, by typing `chown tom comments`, root changes the ownership over to `tom`. Then `tom` can work with this file and even change its permissions to something he prefers.

Hangin' out in groups

Groups are more interesting to work with than owners. You use groups to allow the root user to assign to multiple users the ability to share certain file-system areas. For example, in many versions of Linux, all users are added to a group named *users* (SuSE does this, for example). Then rather than a long format file listing such as the one shown in earlier in this chapter, you may see the following:

```
total 20
drwx----- 2 dee users 4096 Jul 29 07:48 .
drwxr-xr-x 5 root root 4096 Jul 27 11:57 ..
-rw-r--r-- 1 dee users 24 Jul 27 06:50 .bash_logout
-rw-r--r-- 1 dee users 230 Jul 27 06:50 .bash_profile
-rw-r--r-- 1 dee users 124 Jul 27 06:50 .bashrc
-rw-rw-r-- 1 dee users 0 Jul 29 07:48 lsfile
```

In other distributions (such as Fedora), a unique group is created for every user, which is why the earlier listings showed the owner and group items as identical (dee dee).

Comprehending file types

The first letter in any long format file listing tells you which type of file you're dealing with. In Table 7-3, I list the types you're likely to run into. Some of these may seem familiar after examining the file listing color scheme.

Table 7-3		Linux File Types
Label	Type	Description
-	Regular file	The item is an everyday file, such as a text file or program.
b	Block device	The item is a <i>driver</i> (control program) for a <i>storage medium</i> , such as a hard drive or CD-ROM drive.

<i>Label</i>	<i>Type</i>	<i>Description</i>
c	Character device	The item is a <i>driver</i> (control program) for a piece of hardware that transmits data, such as a modem.
d	Directory	The item is a container for files, also referred to as a <i>folder</i> in some operating systems' lingo.
l	Link	The item is a link, as discussed earlier in this chapter.

In addition to those mentioned in Table 7-3, you find lots more different file types out there in the Linux world. By types, I'm not referring to extensions, such as `.exe` or `.doc`. Linux sees everything within its filesystem — even directories and hardware like your monitor — as “files.” As a result, assigning a type to a file is merely a Linux machine's way of keeping track of what's what.

The main thing Windows users in particular want to know when they move to Linux is how to recognize programs. Instead of looking for files with particular extensions (like `.exe`), programs have (or need to have) an *executable permission* set so the system knows they are allowed to run. See Chapter 8 for how to do this in Nautilus.

On the command line, try out the `file` command, such as `file Desktop`, to find more about what a particular file contains.

Navigating the filesystem in Fedora

Fedora's filesystem navigation is all handled through Nautilus (see Chapter 8) if you're using its GNOME default. Some things are made pretty easy for you by default:

- ✓ When you insert a data CD-ROM or DVD-ROM, an icon appears along the left side of your desktop. Double-click this icon to open the item's contents. You can later remove it from the system by right-clicking the icon and choosing Eject from the subsequent shortcut menu.
- ✓ When you plug in a USB storage device such as a keychain, Fedora automatically adds an icon onto your desktop for that item, which you can double-click to open. After you're finished working with the USB device, close all the windows that were using it, right-click its icon on the desktop, and choose Unmount. Now it's safe to remove the keychain. It is vital that you unmount such a device before just removing it; you might lose data otherwise.



- ✓ When you insert a music CD, Fedora opens a music player and begins playing it.
 - ✓ When you insert a video DVD, Fedora attempts to load and play the DVD, but you may need to add more software to make this feature work due to legal reasons. See Chapter 14 for more information.
- To change how Fedora handles all these items, choose System⇨Preferences⇨Hardware⇨Removable Drives and Media. If you shut off auto-loading for something, you can access a device's contents from any Nautilus window by choosing Places⇨Computer and double-clicking the appropriate icon in there. This action adds an icon onto your desktop as well.
- ✓ When you insert a blank CD or DVD writeable or rewriteable, Fedora opens Nautilus to the CD/DVD Creator window. (You can find out more about burning in Chapters 3 and 8.)
 - ✓ When it comes to floppies, the computer can't reliably detect that a floppy was put into the drive. So, put the floppy into the drive, choose Places⇨Computer, and then double-click Floppy Drive to access the disk's contents and add a floppy icon on your desktop. To remove the floppy, right-click the floppy icon on your desktop and then choose Unmount Volume. Wait until the floppy drive light is off before removing the floppy.
 - ✓ If you want to access a network drive from Windows or that someone has set up on another Linux computer, and the ability to share has already been set up, choose Places⇨Network Servers. This action lets you browse through your network to the computers that are offering files.
 - ✓ If you want to access part of your Linux filesystem, choose Places⇨Computer and then double-click the Filesystem icon.
 - ✓ Access something on one of your hard drives that isn't part of your Linux installation (maybe a Windows drive).

For the last option listed, follow these steps:

- 1. Open a command line terminal.**

See Chapter 6 if you're not sure how to do so.

- 2. Type `su -` and press Enter to become the root (administrative) user.**

You will be asked to enter the root password. Do so when prompted and press Enter.

- 3. Type `fdisk -l` to see all your hard drives and partitions.**

If you're looking for a Windows partition, you can find it right here. Ignore the items that have text similar to Windows 95 Ext'd; they're not really a data partition. Anything that has NTFS or FAT32 or VFAT is a Windows partition and is what you're interested in. If you're looking for a Linux partition (maybe you have more than one distribution installed on your system), you want the ones that have the word *Linux* in their description and not the word *swap*.

4. When you think you know what partition you want to try, type `mount -t type /dev/partition /mnt` to add it to your filesystem, where *type* is the filesystem type, as shown in Table 7-4, and *partition* is the letter and number combination you saw in `fdisk` (such as `hde2`).

So, for example, after adding NTFS support, you might type `mount -t ntfs /dev/hde2 /mnt`. When you do so, you can find all of that partition's contents under the `/mnt` directory.

Table 7-4	Common Filesystem Types
Type	Description
ext2	Older Linux filesystem type
ext3	Newer Linux filesystem type, used by default in Fedora and Red Hat, among others
ntfs	Newer Windows filesystem type, started with Windows NT and commonly used in Windows XP
reiserfs	Another newer Linux filesystem type, used by default in Mandrake and SuSE, among others
vfat	Older Windows filesystem type, such as FAT32 and MS-DOS

5. Browse as you need to.

See Chapter 8 for how to browse your filesystem by using Nautilus.

Partitions versus Directories

One very important (and geeky) thing to understand about the Linux filesystem is that it may not all be on one single hard drive or hard drive partition, and yet you don't have to keep track of what drive or partition it's on like you do in other operating systems. In the Microsoft Windows world, if you use separate hard drives or partitions, you have a specific letter designation for each one. The primary hard drive is C, the next is D, and so on. Under Linux, each of these drives and partitions quietly blends together.

If you partitioned your hard drives on your own during or before the installation process, you know that you needed to specify a *mount point* for each partition — which is like an empty spot in a puzzle, where the outside partition or media can be plugged into the rest of the filesystem. In the case of a hard drive partition that's part of your primary main filesystem structure and added at boot time, there's no need for anyone to notice which directory is on which partition.

However, in other cases, it will be more obvious that what you're working with is just a docking ship in the night, not part of the main island, as it were. At these points, if you're working on the command line, you will have to know where to find, say, the CD-ROM that you inserted into the CD-ROM drive (if you're using the graphical desktop, you just point and click your way there as discussed in Chapter 8). Hint: Look in `/media` in your filesystem.

The times you do need to know how Linux sees the hardware are when you're trying to add new hard drives, install the machine while not using automatic partitioning, or access temporary media. Table 7-5 shows you some of the common shortcuts used for your device drivers, for those times you need to access disks from the command line.

Table 7-5 Common Drive Designations	
<i>Designation</i>	<i>Description</i>
<code>/dev/cdrom</code>	CD-ROM drive; if you have more than one, you may have <code>/dev/cdrom1</code> and so on, and you also may see <code>/dev/cdwriter</code> , <code>/dev/dvd</code> , or even a hard drive designation such as <code>/dev/hda</code>
<code>/dev/fd0</code>	Floppy drive 1
<code>/dev/fd1</code>	Floppy drive 2
<code>/dev/hda</code>	First IDE hard drive
<code>/dev/hda1</code>	First IDE hard drive, first primary or extended partition
<code>/dev/hda2</code>	First IDE hard drive, second primary or extended partition
<code>/dev/hdb</code>	Second IDE hard drive
<code>/dev/hdb1</code>	Second IDE hard drive, first primary or extended partition
<code>/dev/hdb2</code>	Second IDE hard drive, second primary or extended partition
<code>/dev/sda</code>	First SCSI hard drive
<code>/dev/sda1</code>	First SCSI hard drive, first primary or extended partition, and often also a small USB drive like a thumbnail/keyring storage device

You probably see a pattern by now. A hard drive has a three-letter designation:

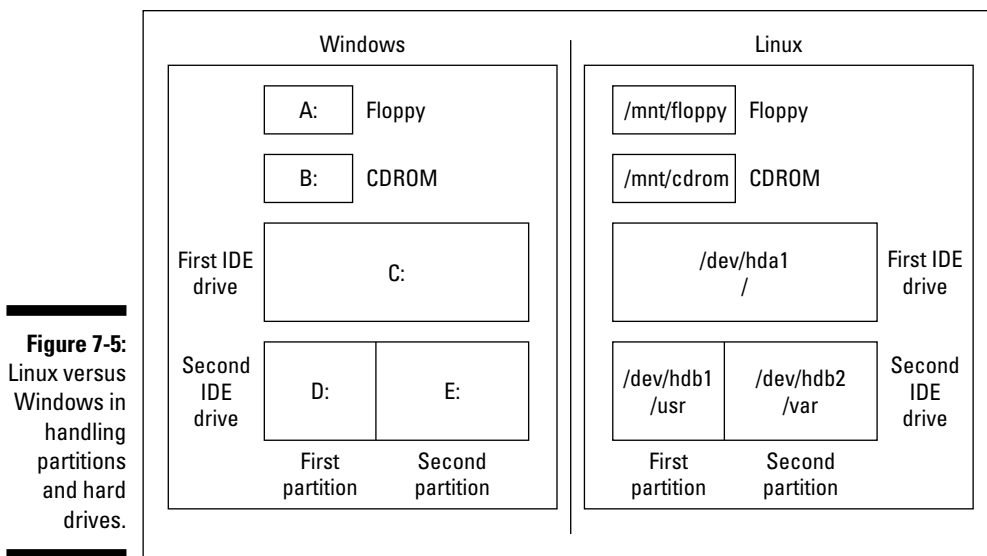
- ✓ An IDE drive's name starts with `/dev/hd`; the first drive of this type is `a`, the second is `b`, and so on. The third IDE drive looks like this: `/dev/hdc`. If you're using something like a Promise controller that bypasses the main IDE controllers (I know, this is technical stuff), your first IDE hard drive

will, in fact, show up as `/dev/hde`. I know this one from personal experience!

- ✓ A SCSI drive's name starts with `/dev/sd`; the first drive of this type is also `a`, the second is `b`, and so on. The fourth SCSI drive looks like this: `/dev/sdd`. USB drives are also seen as SCSI devices. A little item like a USB keychain is often seen as `/dev/sda1` by your system if you don't have any permanent SCSI drives attached. (If you aren't sure if you have IDE or SCSI, just keep in mind that you usually have to go out of your way to get SCSI. If you didn't, you probably have IDE.)

The number that follows the three-letter name represents the partition you're referring to. I cover partitioning your hard drive in Chapter 2.

In Figure 7-5, I break down this concept, hopefully making it a bit more accessible. In this case, the user created three partitions for Linux. The first IDE drive is a single partition, allocated for the `root` partition. The second IDE drive is broken into two partitions. The first was given `/usr`; and the second, `/var`.



In case you're enjoying mucking around on the command line, here's a short primer on some other commands you'll find useful when it comes to wandering through your filesystem. To move, you can use `cd` to change from one directory to another. For example, to go from `/home/sanji` to `/home/sanji/Documents`, you could type one of the following:

```
cd /home/sanji/Documents
cd Documents
```

Either of those commands will get you there. Then to go back to `/home/sanji` from `/home/sanji/Documents`, you could type one of the following:

```
cd /home/sanji
cd ..
```

Again, either of these commands works for the purpose. In both cases, the first commands are referred to as *absolute* navigation (“I want to go to exactly this address”) and the other is called *relative* (“I want to go three houses down from here”). The “..” backs you up one directory in the tree. If you get lost and can’t figure out which branch of the tree you’re on, type **pwd** to find out what directory you’re in right now.

Now some quick commands to work with files. To work with their contents, see Chapter 11. However, just to copy files at the command line, use the `cp` command such as the following to copy `file1` to a new file named `file2`:

```
cp file1 file2
```

If you want to rename a file, use the `mv` (move) command in the same format, or use this same command to move a file from one location to another. To create a directory, use the `mkdir` command such as the following to create a `Documents` directory in the current directory:

```
mkdir Documents
```

To delete files, use the `rm` (remove) command, such as this format to delete `file1` from the directory `/home/keiko`:

```
rm /home/keiko/file1
```

Finally, to remove a directory, with the special `rmdir` command, first you have to delete all the files inside of it, and then you can delete it. To remove all the files in the directory `/home/keiko/notes` and delete it at the same time (make sure you really want to do this!), you can type either

```
rm -rf /home/keiko/notes
```

or you can type the following series of commands:

```
cd /home/keiko/notes
rm *
cd ..
rmdir notes
```

This should be enough to get you started. Happy exploring!

Chapter 8

Using the Filesystem in GNOME and Nautilus

In This Chapter

- ▶ Sailing the filesystem with Nautilus
 - ▶ Working with files in Nautilus
 - ▶ Using CDs and other removable media
-

There is no need to do any housework at all. After the first four years, the dirt doesn't get any worse.

— *Quentin Crisp*

In Chapter 5, we took a look at GNOME and how to use the desktop, but one skill was deliberately skipped: handling files. Chapter 7 addresses how to work with the filesystem on the command line and where to find things. This chapter focuses on pointing and clicking your way through your directories and file manipulation with ease.

Sailin' through Folders

In the GNOME desktop environment, the file manager is *Nautilus*. Moving through the filesystem in Nautilus involves a couple of different skills than using the command line. You may know exactly where in the filesystem you want to go, but not what to click to get there. In the first place, you need to know where you want to begin. This decision isn't as difficult as it sounds.

You have three options for navigating your filesystem with Nautilus:

- ✓ **Double-click your desktop's Home icon:** Double-clicking the Home icon (Figure 8-1) is the same as going to Places⇨Home Folder (see Chapter 5). Note that the window is in a single pane (the block of white). Along the top is a menu bar, and at the lower left is a navigation drop-down list box

that lets you quickly move to parent and “parent of parent” folders. Notice that this window includes a Places menu.

- ✓ **Select Applications→System Tools→File Browser:** Choosing this option (Figure 8-2) gives you a two-pane window with navigation options from Places on the left. (Click the Places drop-down list box to see what other options are available in the left pane, or click the X in the upper portion of the left pane to close it.) This file browser window looks more like a Web browser than the previous option. You can return to the Places section at any time by using the drop-down list box.
- ✓ **Select Places:** This menu gives you the options discussed in Chapter 5, letting you immediately open Nautilus to a predefined location, using the format shown in Figure 8-1.

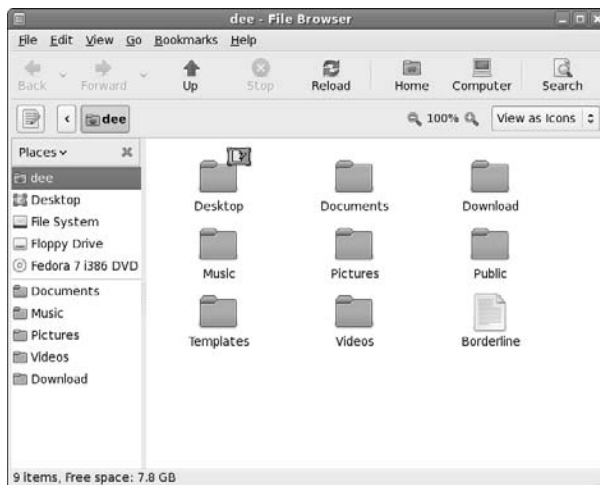
Figure 8-1:

The user dee’s Home folder in Nautilus when double-clicking the Home icon on the desktop.



Figure 8-2:

The user dee’s Home folder in Nautilus when selected from the Places menu.



Regardless of how you opened the window, much of the navigation is the same:

- ✓ To use a file, double-click it. If you are asked for instructions, proceed to the section “Opening files and running programs,” later in this chapter.
- ✓ To open a folder in the layout shown in Figure 8-1, double-click it. A separate window opens with the folder’s contents, leaving the parent folder’s window open as well.
- ✓ To close a window, choose **File→Close** or click the X in the window’s upper-right corner.
- ✓ If you are in the view shown in Figure 8-2, you can navigate the filesystem more easily by clicking **Places** in the left pane and selecting **Tree** from the drop-down list box. Then click the right arrows to expand folders to see what other folders are inside. (Home Folder refers to your account’s files and File System refers to the root folder discussed in Chapter 7.) Click down-facing arrows to collapse sections of the tree and clean up the view. Click the folder itself on the left to see its contents on the right. Double-click folders on the right to see their contents replace the current view’s contents. Notice that there is no **Places** menu along the top; instead, it’s an option in the left pane.

If you’re using the layout shown in Figure 8-1 and find that you’ve got way too many directories open, choose one of the following commands:

- ✓ **File→Close Parent Folders:** Closes all the folders used to get to this directory.
- ✓ **File→Close All Folders:** Closes all the file browser folders you have open.
- ✓ **File→Close:** Closes just this folder.

Setting how much you see in a folder


Are you getting too much information about each file or not enough? Experiment with the **View** menu. The major listing options under the **View** menu include the following commands:

- ✓ **View As Icons:** This default option shows all files as icons with names beneath them; refer to Figure 8-1.
- ✓ **View As List:** This command gives you a list of information about the file, as well as its icon and name (see Figure 8-3). Note that there are still small icons available along the left of the listing.

Along with these **View** menu commands, you can also use

- ✓ **View→Reload:** If you just created a file and don’t see it in the folder, select this option.
- ✓ **View→Show Hidden Files:** All files in Linux whose names start with a dot (for example, `.profile`) are *hidden files*, meaning that you have to explicitly tell Linux you want to see them.

Figure 8-3:
Nautilus
with the
View As List
option in
dee's home
directory.



Along with these commands, you can choose the View menu's Zoom In, Zoom Out, and Normal Size options if you need to change the visual size of the folder's contents. Choose the Reset View to Defaults option if you want to start over from scratch. For more advanced settings for the View As List mode, try View ⇨ Visible Columns.

Creating files and folders

To create a new file or folder in Nautilus:

- 1. Browse to the folder you want to create the item in.**

You need to have the folder open and be inside its window before proceeding.

- 2. Right-click on an empty spot inside that folder.**

Make sure that you didn't highlight anything in the folder. A context menu appears.

- 3. Depending on what you want to accomplish, choose either Create Folder or Create Document.**

- **The Create Folder command** creates a new, empty folder in this directory.
- **The Create Document command** opens the submenu that lets you choose to create either a document from a Template (top section, if you already have templates created) or an Empty File.

The folder or file appears, highlighted and with the name open for editing.

4. Type the name you want to assign to this folder or file and then press Enter.

The name is assigned to the folder or file.

Opening files and running programs

To open a file or run a program in Nautilus, double-click it. That's it! Well, not entirely. You can also right-click a file to open the context menu and then choose one of the following from the context menu:

- ✓ **Open with *program*:** If the file is associated with a particular program already, choosing this option is the same as double-clicking the file. More than one of these options may be listed, if your system knows that you have more than one program installed that can do the job — the top item with the folder icon next to it is the default. Choose whichever option you prefer.
- ✓ **Open with Other Application:** Choosing this option opens a submenu that contains a list of installed programs the system is aware of. You can select one of these programs if you find what you need in there by clicking Use a Custom Command. From here, you can
 - Type in the full path to the application (the full directory path, like `/usr/bin/konqueror`).
 - Click the Browse button to open a file browser and surf to the program you want to use.

After you select the program to open the file with, click Open to proceed.

Copying and moving files

You can copy and move items by using two different methods in Nautilus. The first is by using the usual method you're probably familiar with from Windows or the Mac OS — clicking and grabbing a file or folder and then dragging it where you want to go. You can click and drag between folder windows and into folder icons. Following are some handy tips:

- ✓ Hold down the Ctrl key while you drag if you want to make a copy rather than just move the file.
- ✓ If you want to drag the file into a folder icon, make sure that the folder icon is highlighted before you release the mouse.
- ✓ If you want to drag the file into a folder window but not into one of the folder icons, make sure that the folder icons are *not* highlighted before you release the mouse.

The other method for copying and moving files and directories involves the following steps:

1. **Right-click the file or directory you want to copy or move.**
2. **From the context menu that appears, choose Cut if you want to move the file or Copy if you want to copy it.**

The file doesn't disappear after you make your selection if you're trying to move it.

3. **Navigate into the folder you want to move the file into.**

You are ready when you have that folder's window open.

4. **Do one of the following:**

- If you want to move or copy the file into a folder window, right-click inside that window (just make sure that nothing is highlighted) and then choose Paste.
- If you want to move or copy the file into a folder icon rather than an open folder window, right-click the icon and then choose Paste Into Folder.

Deleting files and folders

To use Nautilus to delete either a file or a directory from the filesystem, follow these steps:

1. **Browse to the file or directory's location.**
2. **Select the file(s) or folder(s).**

You can select the file or folder by doing one of the following:

- Clicking the file or folder to highlight it.
- Clicking the first item, holding Shift down, and then clicking the last item in the group to choose them all. (These items all must be lined up in a row.)
- Clicking the first item, holding Ctrl down, and then clicking each individual item you want to select.
- Clicking and dragging your mouse button so that you make a box that contains all the items you want to select.

3. **Press the Delete key or right-click and select Move To Trash.**

The file or folder vanishes from view. Keep in mind that if you are deleting a folder and the folder contains other files or folders, they go to into the trash right along with it.



These deleted items are in fact sitting in the Trash can, so they aren't really deleted from your hard drive. You can move these files out of the trash and back into your filesystem if you want to.

Taking out the trash

You can permanently delete the contents of your Trash folder by following these steps:

- 1. Right-click the Trash icon.**

This action opens a shortcut menu with options listed.

- 2. Choose Empty Trash from the shortcut menu.**

A confirmation dialog box opens.

- 3. Click Empty in the confirmation dialog box to delete the contents of the Trash folder.**

Linux permanently removes the items in the Trash folder. You can tell whether the Trash folder is empty by looking at the Trash waste can. If it's empty, it contains no files. If you see papers in it, it contains trash.

Viewing and changing permissions

In Chapter 7, you see that every file in Linux (and other forms of UNIX) has a set of *permissions* that govern who is allowed to view it, run it, delete it, and so on. These permissions are used to make sure that people can't mess with the system's or each other's files, so they're pretty important. In this section, I focus on how to work with files in the GUI.

To view and change a file or directory's permissions in Nautilus, you need to follow these steps:

- 1. Browse to the file or directory's location.**

Opening the folder window that contains this item is sufficient.

- 2. Right-click the file or directory.**

A context menu appears.


- 3. Choose Properties from the context menu.**

The Properties dialog box appears with the Basic tab open.

- 4. Click the Permissions tab.**

The Permissions portion of the Properties dialog box appears, as shown in Figure 8-4.

Figure 8-4:
The Nautilus
file man-
ager's
Properties
dialog box
with the
Permis-
sions tab
displayed.



5. Set the new permissions and ownerships.

See Chapter 7 for a breakdown of what all this stuff means. Notice that the GUI has the user, group, and “others” broken down into separate drop-down list boxes. Unless you are familiar with SELinux, leave the SELinux Context option alone.

6. Click Close to close the dialog box.

The file’s permissions are now changed.

Using CDs and Other Removable Media

As discussed in Chapter 7, when it comes to removable media (floppy disks, CD-ROMs, DVD-ROMs, USB thumb drives, and so on), you end up having to look around for them if you’re working on the command line. The GUI makes finding the media you’ve attached a much easier process.

Typically speaking, when you insert a new removable item (say a data CD-ROM), an icon appears on the left of your Desktop with the name of the device assigned to it. You can double-click the icon to see the media’s contents. If an Eject option is available (for example, for the CD-ROM), you can right-click the icon and choose Eject from the context menu. In cases where Eject is not available, use Unmount before removing the item (such as a USB thumb drive) from the computer. This action ensures that all the data you saved actually is written, because Linux doesn’t always put the data on the drive right away.

In most cases, when you insert a music CD or DVD, a music player (see Chapter 13) opens to play the music. See Chapter 13 for how to add support

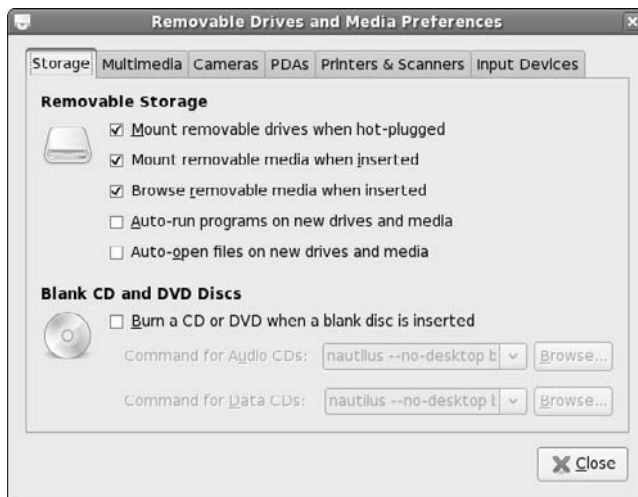
for even more types of music files. When it comes to video CDs and DVDs, especially video DVDs, you often have to install extra software due to a number of issues discussed in Chapter 14.

When it comes to blank CD or DVD writeables or rewriteables, inserting these into their drives causes a Nautilus window to open to its CD/DVD Creator window (see Chapter 13 for how to burn your own CDs and DVDs).

However, the computer can't reliably detect that a floppy was put into the drive, so you have to manually tell it to look after you have inserted the disk. So, put the floppy into the drive, choose **Places**⇨**Computer**, and then double-click **Floppy Drive** to access the disk's contents and add a floppy icon on your desktop. When you're finished working with the floppy disk, you can remove it by right-clicking the floppy icon on your desktop and then choosing **Unmount Volume**. Wait until the floppy drive light turns off before removing the disk.

If you want to change how Fedora handles such media, choose **System**⇨**Preferences**⇨**Hardware**⇨**Removable Drives and Media** to open the **Removable Drives and Media Preferences** dialog box (Figure 8-5). Most of the settings applicable to this section are in the **Storage** tab.

Figure 8-5:
The
Removable
Drives and
Media
Preferences
dialog box
with the
Storage tab
displayed.



Some of the options are a bit hard to follow, so here's an explanation of some of the terms used:

- ✓ **Mount** means that the drive should be added to the filesystem so you can use it, and an icon appears on your desktop showing you that it has been added.
- ✓ **Hot-plugged** devices can be plugged in or unplugged from a system while it's running, like plugging in a USB thumb drive.

- ✓ Browse refers to mounting but also opening a Nautilus window with the media's contents.
- ✓ Auto-run is as you would expect from the Windows world, where the system will try to identify the program it should run when the media is inserted, and start that program automatically.
- ✓ Auto-open is similar to auto-run, except that files will be opened to display their contents.

Clicking the Multimedia tab (Figure 8-6) allows you to specifically deal with the settings for music and video. The settings available here are

Figure 8-6:
The
Removable
Drives and
Media
Preferences
dialog box
with the
Multimedia
tab
displayed.



- ✓ In the Audio CD Discs section, you can set whether or not an audio CD should play automatically when placed in the drive. Beneath that, the Command text box tells the system what program to play when the disc is inserted. See Chapter 13 for a suggestion of a better player than the default, `totem` (this suggested player is actually shown in Figure 8-6 rather than the default).
- ✓ In the Video DVD Discs section, the settings are similar. Here you can also set what happens when the disc is inserted, and what command should be run.
- ✓ In the Portable Music Players section, if you have an MP3 player, you can set the system to play the files automatically when the player is connected. If the box is checked, you can also set the command to use in the Command text box. I recommend `amarok %d` (see Chapter 13) unless you need iTunes (see Chapter 15).

If you have other devices you want to set up, feel free to explore the other tabs. When you're finished with this dialog box, click Close to save your changes.

Accessing Windows Drives on This Computer

If you're dual-booting between Windows and Linux, you will have times when you want to access information on your Windows drive from Linux (unfortunately Windows cannot access the information on Linux drives). To access the files on your Windows drive from Linux when they're both on the same computer, do the following:

1. **Open a command-line terminal.**

See Chapter 6 if you're not sure how to do so.

2. **Type `su -` and press Enter to become the root (administrative) user.**

You are asked to enter the root password. Do so when prompted and press Enter.

3. **Type `fdisk -l` to see all your hard drives and partitions.**

If you're looking for a Windows partition, you can find it right here. Ignore the items that have text similar to Windows 95 Ext'd; they're not really a data partition. Anything that has NTFS or FAT32 or VFAT is a Windows partition and is what you're interested in. If you're looking for a Linux partition (maybe you have more than one distribution installed on your system), you want the ones that have the word *Linux* in their description and not the word *swap*.

4. **If you need to access an NTFS partition, see Chapter 16 for how to add NTFS support to Fedora.**

The packages you want all start with `ntfs`. There's a collection of them: Go ahead and install them all, or at least `ntfs-3g`, `ntfs-config`, and `ntfsprogs`.

5. **When you think you know what partition you want to try, type `mount -t type /dev/partition /mnt` to add it to your filesystem, where *type* is the filesystem type (`vfat` or `ntfs`), as shown in Table 8-1, and *partition* is the letter and number combination you saw in `fdisk` (such as `hde2`).**

So, for example, after adding NTFS support, you might type `mount -t ntfs /dev/hde2 /mnt`. When you do so, you can find all of that partition's contents under the `/mnt` directory.

Table 8-1	Common Filesystem Types
Type	Description
ext2	“Older” Linux filesystem type
ext3	“Newer” Linux filesystem type, used by default in Fedora and Red Hat, among others
ntfs	“Newer” Windows filesystem type, started with Windows NT and commonly used in Windows XP and Vista
reiserfs	Another “newer” Linux filesystem type, used by default in Mandrake and SuSE, among others
vfat	“Older” Windows filesystem type, such as FAT32 and MS-DOS

6. Browse as you need to.

You may find that you have read-only access to your NTFS partition and can't change anything on it.

7. If you want to be able to write to the NTFS drive, go to Applications⇨System⇨NTFS Configuration Tool, enter your root password if necessary, click the Enable Write Support for Internal (if the drive is in your computer) or External (if the drive is connected from outside) Device check box, and click OK.

8. When you're finished with the drive, type `umount /mnt` to remove the partition's contents from your filesystem.

Accessing Network Drives

Sometimes the drive you want to access is on another computer. If the drive on the other computer is set up to be accessible through your network, select Places⇨Network. A Nautilus folder opens with your network navigation options available. If you want to access a Windows computer, double-click Windows Network. The Windows computers on your network that have drives available for browsing will be displayed. From there, you can browse to where you want to go!

Formatting Disks

A floppy disk, USB keychain, and any other small(ish) storage device often comes as a blank slate *or* formatted for Windows or Macintosh use (most often these days they come formatted for Windows). If the item is a blank slate, no computer can use it for anything. Many come by default formatted for Windows, which you can use in Linux with no problem — or you can change its formatting into a Linux setup. (The handy thing about leaving it as a Windows disk is that you can then use it to share things with Windows users.)

To format a floppy, place the floppy into your floppy drive and choose Places⇨Computer. In the Computer dialog box, right-click the Floppy Drive icon and choose Format from the context menu.

Typically, the defaults are what you want to choose. If you want to format the disks for Linux, use the Linux Native (ext2) format. To share disks with Windows users, choose DOS (FAT). Another setting you may want to choose is Thorough rather than Quick.



This same technique works for many types of removable media. Locate the appropriate icon in the Computer dialog box and right-click it just as you would for a floppy. Many devices like USB thumb drives are sold pre-formatted for DOS (FAT). It is fine to leave them in this form.

Burning Data CDs and DVDs

A CD and DVD burner is a great way to make data backups (especially if it's a rewriteable), save and share your digital photos, put together multimedia memento scrapbooks, and more. Many tools are available that allow you to burn CDs under Linux, but you can also burn CDs and DVDs by using the Nautilus. Just right-click the file you want to burn, and choose the proper options from the context menu.

However, you may want to use K3b if you want to get fancy. If you don't have it installed yet, see Chapter 16 for how to do so — the specific packages you want are `k3b` and `k3b-extras`. To open it, go to Applications⇨Sound & Video⇨K3b. When you launch this tool, you see the window shown in Figure 8-7.

Figure 8-7:
The K3b CD
and DVD
Kreator.

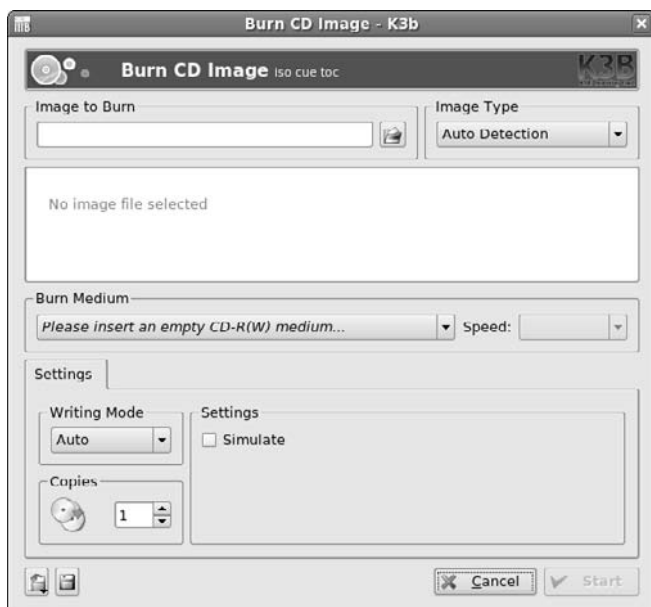


If all you want to do is create a CD or DVD out of an ISO file you downloaded, do the following:

1. **Select either Tools⇒Burn CD Image if this is an ISO file for a CD-ROM, or Tools⇒Burn DVD ISO Image if this is a file for a DVD-ROM.**

Both actions open the Burn CD Image dialog box (Figure 8-8).

Figure 8-8:
The K3b
Burn CD
Image
dialog box.



2. In the Image to Burn section, click the folder button and navigate to the file you want to burn onto CD or DVD.

3. When you find it, select it and click OK.

The filename and information now appear in the dialog box.

4. If you have more than one burner for some reason, make sure that the proper one appears in Burn Medium, and change the selection if necessary.

5. In Speed, if you are burning a CD or DVD of a Linux distribution, lower it to at least 4x. Otherwise, you are usually safe to leave it as is.

If you find that when you burn CDs or DVDs, they often seem to end up as “coasters” (burns that didn’t work), lower the speed you usually use.

6. Click Save User Defaults if you want to use most of these settings for every burn.

7. Click Start to begin your burn.

A dialog box appears with two progress bars. The first shows you the write preparation progress, and then the second shows you the total progress, including burn time.

8. If this is a rewriteable media, you’ll be asked if you want to erase its contents so it can be written over. Answer appropriately.

If you tell K3b to erase the media, the Erasing dialog box appears during that part of the process. When the erasure finishes, the tray opens, ejecting the media, and then automatically closes again. This moment is when the actual burning session starts. The next time the disc ejects, the process is complete.

9. Remove your new CD or DVD from the drive and then click Close to close the burn information dialog box, and Close again to close the main burn dialog box.

10. Select File→Quit to close K3b.

Notice that this program has some other interesting features down in the bottom portion of the main window. Clicking one of the Project buttons will help you create a music (audio) CD, a data DVD, or a data CD from scratch. Making audio CDs is covered in Chapter 13.

Finding Things

A wonderful tool called Beagle gives you the capability to search both file-names and inside files, making it incredibly useful for tracking down all kinds of things on your desktop. This tool isn’t installed by default, so use the instructions in Chapter 16 to install all of the programs with the word

“beagle” in their name. After you’ve done so, you can run this program by going to Places⇨Search (Figure 8-9) — notice that Beagle has now taken the place of the more primitive search tool that was there before. You will find a set of Quick Tips in the main window for how to make the best use of this tool.

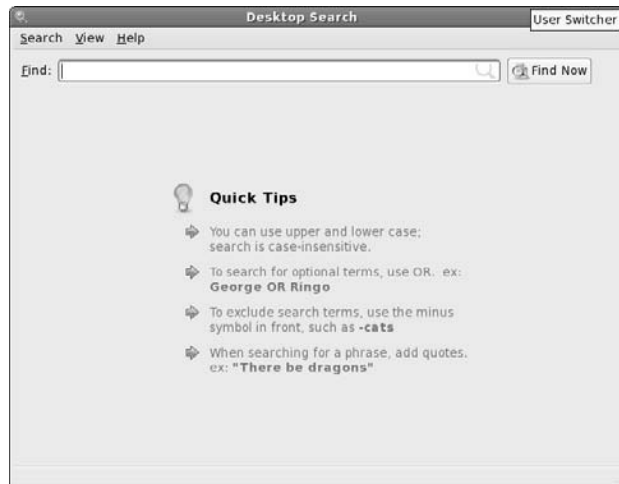


Figure 8-9:
The Desktop
Search tool.

For an example of how useful Beagle is, if you used OpenOffice.org Writer to write notes for a report on the Renaissance, you could type in Renaissance, click Find Now, and everywhere Renaissance appears (Documents, Conversations in IM, even your Web surfing cache) will come up in the search results.




If you see a button appear that says Start Search Service, click it, and the program that runs in the background for your searches will start up. After it’s started, the search you were trying to run completes automatically.

To see all the possible categories you can search (and adjust which ones the system should bother with), choose Search⇨Show Categories. In the sub-menu, you will see the full list of the types of material Beagle can and should access. You can remove an item individually from the list by selecting it, or remove them all by selecting None and then individually choosing the items to add back.

Selecting Search⇨Preferences opens the Search Preferences dialog box (Figure 8-10), where you can configure how Beagle behaves for your particular login.

Figure 8-10:
The Desktop
Search
tool's
Search
Preferences
dialog box.



The tab you'll be most interested in is the Indexing tab. Here, you can tell Beagle to look outside of your home directory (where your personal files are stored) by, under General, clicking Add. Or, you can use the Privacy section to Add items that should not be indexed.

Chapter 9

Connecting to the Internet

In This Chapter

- ▶ Understanding common Internet connection methods
 - ▶ Setting up your Internet connection
 - ▶ Connecting to your Internet Service Provider (ISP)
 - ▶ Understanding enough TCP/IP to be dangerous
-

Every improvement in communication makes the bore more terrible.

— Frank Moore Colby

You may already be connected to the Internet if you're on a machine that's connected to a LAN, and you've configured the networking during installation. To test whether you have a connection, open a Web browser and try to go to an outside Web site (like `www.gnu.org`). If it works, you're up! No need for this chapter. Otherwise, read on.

Internet Connectivity 101

A few words regarding broadband options (DSL and cable) are in order before I begin. The word *broadband* has a technical definition, but I just use it to mean high-speed Internet access. Although a dialup modem can typically transmit information at speeds up to 56 *kilobits* (thousand bits) per second, broadband connections can reach 50 times that speed. Nowadays, the Web contains lots of images and multimedia elements, and enjoying these features through a 56 Kbps dialup modem is similar to drinking a cold glass of water with an eyedropper when you're dying of thirst.

Don't let high-speed providers discourage you from using Linux with their services. Just because they don't support Linux directly doesn't mean that the technology doesn't work with Linux. TCP/IP (the set of traffic rules for the Internet) was developed for the UNIX operating system, from which Linux descends.

If you are dual-booting to Windows, your ISP can help you install your broadband connection, and then you can tinker with getting Linux connected as you have the time and inclination. In the list below, I point you to some sites that can help you configure your broadband connection as you look at the types available:

- ✓ **Cable modems:** When you subscribe to a cable Internet service, the installation technician typically provides you with a special device, called a *cable modem*, along with a standard network adapter (like an *Ethernet card*). The technician then installs the network adapter into your computer and connects the cable modem to it. To find out the exact speeds available, you need to talk to your cable provider. The following Web page provides a document that contains helpful information about various cable providers and Linux:

www.tldp.org/HOWTO/Cable-Modem

- ✓ **Digital Subscriber Line (DSL):** DSL carries data to your telephone jack in a digital format. DSL is popular because it provides a faster connection with lower installation and service costs than ISDN, and it utilizes the existing copper telephone wiring provided by your telephone company. You can even use the same phone line as a regular phone line! A DSL connection requires additional communication hardware, which your Internet Service Provider (ISP) should provide.

Note that several variants of DSL exist:

- ISDN (Integrated Services Digital Network) Digital Subscriber Line (IDSL) — see later in this bulleted list for a definition of ISDN.
- Symmetric Digital Subscriber Line (SDSL).
- Asymmetric Digital Subscriber Line (ADSL).
- Generic DSL (XDSL).

You can find out from your ISP how fast your DSL connection would be, because it varies. For an overview of DSL and Linux, visit the following Web page:

www.tldp.org/HOWTO/DSL-HOWTO/

- ✓ **Satellite modems:** Those who live in particularly remote or low-infrastructure areas might go the route of satellite modems (www.tldp.org/HOWTO/Sat-HOWTO.html). You can consider this service equivalent in many ways to cable, in that you get it typically through the same providers that you get your satellite television signals. However, satellite has some drawbacks. If your satellite television service is unreliable, your Internet service suffers the same fate. The upload and download speeds can also vary dramatically.

Beware of devices posing as modems

Many Linux newbies have become irritated by not being able to communicate with their internal modems. “After all,” the newbie reasons, “the same hardware works when I was running Microsoft Windows.”

Well, here’s the story: Years ago, hardware manufacturers developed a device, called a *software modem*, in an effort to reduce hardware costs. The idea was to trim some responsibility from the modem and relegate these tasks to the operating system. The result was an inexpensive modem that routed signals to proprietary software that operated only under Microsoft Windows. In short, these so-called modems, also known as *WinModems*, aren’t really modems at all, but, rather, are telephone cable interfaces to Windows.

The following list shows methods you can use to determine whether you have a software modem:

- ✓ The model number has a HCF-, HSP-, or HSF- in front of it.
- ✓ The packaging refers to the device as a WinModem or designates that it works with only Windows (though sometimes these labels can be misleading because some manufacturers just don’t bother to list

other operating systems even though their hardware will work with those operating systems).

- ✓ Windows recognizes your modem, but Linux doesn’t.

In short, if you determine that you have a software modem, Linux simply doesn’t work with it. For the adventurous out there, the LinModem project (linmodems.org) has successfully written Linux drivers to work with a few of these software modems. I encourage you to become involved in this type of project, if you’re so moved. These explorers drive the wonderful world of *freedom* software, which Linux is a prominent part of.

The bummer in all this is that, although you probably saved a few bucks by buying a machine with a software modem, you most likely need a real modem to use with Linux (unless you are lucky enough to have a LinModem). Your best bet is to purchase an external modem so that you can leave the software modem in place and have the external one just for Linux (or any other computer you may want to move it to later).

- ✓ **Dialup modems:** The dialup modem is still widely used in areas where broadband is either not available or not affordable. It translates the *digital* signal from the computer into an *analog* signal required for transmission from the wall jack to the telephone company. Because the modem utilizes existing voice telephone service, you don’t need any special setup beyond subscribing to an ISP. However, you can’t use the same phone line for both dialup and a phone conversation at the same time. See the sidebar “Beware of devices posing as modems” for additional concerns if you’re going the dialup route.

- ✓ **Wireless:** The realm of wireless networking is where Linux today may give you the most trouble, depending on what kind of wireless you're using. Four major types of wireless exist, and they're often referred to as a group as *WiFi* (802.11a, 802.11b, and 802.11g, or just A, B, and G, and Bluetooth). A and B are also the best supported in Linux so far, because again, they've been around the longest. G's security is what you should aim for, but it also adds levels of complexity that require more cooperation from the manufacturers of the cards. Many wireless cards today are, in fact, capable of doing more than one of these types, so if you don't own one, get one that supports both B and G.

The best advice I can give here is to do a Web search on the make and model of your wireless card (or the one you're thinking of buying), plus the word *Linux*, and see what you find — and, of course, you can check the hardware compatibility pages for the distributions as well. Because this area is one where technology is quickly advancing, always check the dates on the resources you find and make sure that they're recent. A resource that's just six months old can mean you're getting out-of-date information. Some useful links (some are more advanced than others) for this topic include the following:

- A posting on my blog for Fedora users (apologies for the long URL): http://dee-ann.blog-city.com/wireless_setup_on_an_emachines_m5310_notebook.htm
- The Linux Wireless Networking Quick HOWTO: www.linuxhomenetworking.com/linux-hn/wmp11-linux.htm
- Wireless LAN resources for Linux: www.hpl.hp.com/personal/Jean_Tourrilhes/Linux/
- A method for using Windows wireless drivers in Linux: www.linuxant.com/driverloader
- Network Topology Guide: www.homenethelp.com/network

- ✓ **Integrated Services Digital Network (ISDN):** Not too long ago, ISDN was one of the only residential high-speed options. It appeared when 28 Kbps (half the speed of today's regular dialup modems) was about all that you could milk from the copper strands that connected your telephone to the telephone company. ISDN is still available in some areas and promises a steady 128 Kbps — as long as you're within 3.4 miles of the telephone company's central office. You need two special devices to use ISDN: an ISDN modem (typically provided by the ISDN ISP) and a network adapter (sometimes an Ethernet card). Because ISDN is used so little these days, not many specific sites are devoted to it. The overall modem HOWTO (www.tldp.org/HOWTO/Modem-HOWTO.html) does contain pointers to the latest ISDN information, as well as just being plain interesting to read.

Setting Up the Hardware

Before you get too comfortable in your chair, you must physically check some items. This may require you to do some low-level maneuvers (such as crawling under your desk):

✓ **Cable modem, DSL modem, or any other fancy contraption:** If you have one of these babies, you need to make sure that:

- An Ethernet cable (like a phone cable but the connector is wider) is plugged into your computer's Ethernet card and into the special cable, DSL, satellite, or ISDN modem your ISP installed.
- The fancy modem thing is powered on. (These contraptions often have their own power supplies.)

✓ **External modem:** If you have an external modem (one that's independent of your computer case), you need to verify that

- A cable is securely connected from the modem to the proper port on the computer.
- The modem is powered on. (External modems have their own power supply.)
- One end of a telephone cable is plugged into the wall jack, and the opposite end is plugged into the modem.

✓ **Internal modem:** If you have an internal modem, you need to verify that:

- The modem is *not* a software modem. See the sidebar "Beware of devices posing as modems" if you're not sure what a software modem is or whether you have one.
- One end of a telephone cable is plugged into the wall jack, and the opposite end is plugged into the phone plug on the back of your computer.

✓ **Wireless cards:** If you have a wireless card, you also have to have a wireless *router* to collect and direct your wireless traffic between computers and/or *bridge* to let your wireless card talk to your Ethernet network. How far and through what (walls, floors, and so on) a wireless router or bridge can be from the card depends on the technology you're using. Some handy primers on wireless hardware and setup include

- **How WiFi Works:** <http://computer.howstuffworks.com/wireless-network.htm>
- **Creating a Wireless Network:** www.wi-fi.org/OpenSection/design.asp?TID=2

Okay, now you can climb back into your chair.

Selecting an Internet Service Provider (ISP)

Because of the meteoric rise in the popularity of Linux, many ISPs are training their support staff in the ways of Linux. If you already have a dialup service, give one of them a call to let them know of your Linux pursuits. Chances are, that person already has information pertinent to Linux subscribers and can provide you with that information. If you're shopping for a new ISP, this section offers some practical selection advice.



Some ISPs provide their own proprietary software that you must install on your PC to connect to the Internet. The software they provide is likely to run only on Windows. Several free dialup services don't work with Linux because of this fact. The proprietary software meddles with the operating system to ensure that banner advertising isn't hidden or that you stay dialed in for only a specified duration.

If you're shopping around for an ISP, consider these questions:

- ✓ **Can you get local dialup numbers across the country and around the world?** If you travel often and need Internet access from different cities, this service is a handy money (and headache) saver.
- ✓ **Does it provide technical support for Linux?** If you're planning on running Linux, this consideration is an important one. It doesn't have to be a deal breaker, but it is handy.
- ✓ **Can you get a recommendation?** Ask a friend. An ISP's best friend is an endorsement from a satisfied customer.
- ✓ **Will you have trouble dialing in?** Although a subscriber-to-line ratio of 7-to-1 (an average of seven or fewer subscribers per line) or better isn't an entirely accurate measure of how often you'll get busy signals, it's probably the easiest measure for consumers. This is primarily a concern for dialup users.

Getting Information You Need from Your ISP

Most reputable ISPs provide you with a customer information sheet after you sign up for its services. This sheet should include the following information at the minimum:

- ✓ Local telephone dial-in numbers for modem users
- ✓ User login name
- ✓ User login password
- ✓ E-mail address
- ✓ E-mail outbound host or Simple Mail Transport Protocol (SMTP) server
- ✓ E-mail inbound host or Post Office Protocol (POP) server
- ✓ News host
- ✓ Whether your computer address will be
 - **Static:** You always have the same IP address. Static addresses are more common for servers than for desktops. If you're using static, your ISP needs to give you your nameserver IP addresses, your gateway address, the network address, and the netmask. This information should all be on the customer information sheet.
 - **Dynamic:** Your IP address changes each time you connect, or at regular intervals. On most desktops, you just don't always need to have the same address, and it makes life easier on ISPs. It also makes life easier for computers like laptops, which get moved around often from network to network. If you're using dynamic addressing, you're told that you need to connect by using DHCP.

With this information, you can establish an Internet connection by using your Linux system.

Configuring Your Connection

Okay, so you've got your network hardware and your ISP information, now what? The answer is to configure your Fedora system to jump on the Internet!

The configuration process is slightly different depending on the hardware you're using. This section covers the big three methods for accessing the Internet:

- ✓ Using a modem to dial to your ISP
- ✓ Using an Ethernet card to connect to a DSL or cable modem
- ✓ Using a wireless card to connect to a wireless access point

Each situation requires a few slightly different items to configure, but you can perform all of them by using the Fedora Network Configuration tool.

Dialin' out

Using a modem to dial a connection to your ISP is a two-step process. First, you must configure your modem with the settings from your ISP. After that, you must manually start and stop the Internet connection (you probably don't want your modem tying up your phone line all day long). First, work on setting up the modem.

Setting up the modem

With your ISP information in hand and a glowing monitor in front of you, follow these steps to configure your dialup connection to the Internet in Fedora:

1. From the GUI desktop, choose **System⇨Administration⇨Network**.
2. Enter your root password in the dialog box, if necessary.

If you're logged in as a regular user (not root), you're prompted to enter the root password.

The Network Configuration dialog box opens (see Figure 9-1).



Figure 9-1:
The
Network
Con-
figu-
ration
dialog box.

3. In the Network Configuration dialog box, click the **New** button.

The Select Device Type dialog box appears, as shown in Figure 9-2.


Figure 9-2:
The Select
Device Type
dialog box in
the Network
Configura-
tion tool.



4. To set up a dialup modem, select **Modem Connection** from the **Device Type** list box and then click the **Forward** button.

After you click Forward, the tool probes for your modem. If your modem can't be found (I'll follow through the most difficult option), the Select Modem dialog box appears, as shown in Figure 9-3. I assume that you're staring this dialog box in the face. If it doesn't appear, proceed through your configuration until it matches these steps again.

Figure 9-3:
The Select
Modem
dialog box in
the Network
Configura-
tion tool.



5. Select the proper device from the Modem Device drop-down list.

If you know your modem is on COM1 or COM2 in Windows, translating this techno-talk to Linux isn't hard: You just need to subtract 1. So, it's ttyS0 for COM1 and ttyS1 for COM2. If you're using a USB mouse (it has a flat, rectangular connector), your modem is probably on COM1. If you're not using a USB mouse, then your modem is probably on COM2.

6. Select your modem's speed from the Baud Rate drop-down list.

If you were told to use a specific speed, choose that speed here. Otherwise, choose the highest number available. The key is to try the highest setting and, if that doesn't work, start choosing slower and slower ones until you're able to properly make the connection.

Leave the Flow Control setting alone. Only change this setting if you're having a hard time getting your modem connection to work, or you're told by your ISP that you specifically need to select either Software Flow Control or No (None) Flow Control.

7. Turn on your modem volume so that you can hear what it's doing while testing it.

You'll want to come back later and shut off the volume. (Well, most people do; some like to leave it on. I think the 56.6 Kbps modems sound like an alien game of Pong while they're connecting). It's up to you how loud you want to set the sound, as long as you can hear it.

Leave Touch Tone Dialing activated unless you live in an area that only supports rotary phones.

8. Click Forward to proceed.

The Select Provider dialog box appears, as shown in Figure 9-4.



Figure 9-4:
The Select
Provider
dialog box in
the Network
Configura-
tion tool.



9. Enter your ISP information and then click Forward.

The default list of countries that your ISP may be from is pretty limited. If you live in one of these regions, select your country, and then the ISP — at which point the dialing entries are completed for you — and then skip to Step 16. Otherwise, proceed to the next step.

10. In the Prefix text box, enter any dialing prefix you have to dial before the actual phone number.

This prefix may be a long distance item, such as 1 for North American long distance calls, or a 9 to dial out of your building — or even 91 if you need to dial 9 to get out and you're dialing long distance. If you don't need a prefix, leave this box blank.

11. In the Area Code text box, enter the area code your modem needs to dial if it needs to dial one.

In some places, you have to dial the area code even if you're dialing a local number. If you don't have to use an area code, leave this box blank.

12. In the Phone Number text box, enter the phone number (minus prefix and area code) your modem needs to dial.

Dashes aren't necessary.

13. In the Provider Name text box, type the name of your ISP.

You just use this name as a description for this dialup entry.

14. In the Login Name text box, enter your dial-in account login name.**15. In the Password text box, enter your dial-in account login password.****16. Click Forward.**

The IP Settings dialog box appears, as shown in Figure 9-5.

Figure 9-5:
The IP
Settings
dialog box in
the Network
Configura-
tion tool.



17. In the IP Settings dialog box, select the appropriate radio button for whether your modem needs to obtain its IP addressing information from the ISP (*DHCP*, as discussed in the section “Getting Information You Need from Your ISP,” earlier in this chapter) or whether you need to tell it what its IP data is (*static*).

Your ISP should have given you which of these you require. DHCP is the default. If you choose Static, the lower portion of the dialog box becomes available.

18. If you selected DHCP, and your ISP told you to automatically get your nameserver settings from its equipment, leave the Automatically Obtain DNS Information from Provider box checked. Otherwise, uncheck it.
19. If you selected static, fill in the Address, Subnet Mask, and Default Gateway Address as given to you by your ISP.
20. Click Forward to proceed to the summary dialog box, and then click Apply to add the modem you have just set up to your main Network Configuration dialog box.

You now have a modem entry (look in the Type column) in your Network Configuration tool.

Connecting to (and disconnecting from) the Internet

To connect to the Internet from the Fedora Network Configuration dialog box, highlight the modem entry you just created (see preceding section) and click the Activate button. With any luck, your modem springs to life with some beeps and then buzzes along with a dialog box indicating that a connection is being made. After a successful connection, the Network Configuration dialog box appears again; this time, the Status column of the modem device reads Active. Congratulations — you’re now connected to the Internet!

To disconnect, click the Deactivate button in the Network Configuration dialog box. Your settings remain; the next time you want to connect, just open this dialog box and click Activate. If all is not well, skip to the “It’s All Fun and Games Until Something Doesn’t Work” section to find out how to get help.

Playing with cards

With the increase of broadband connections, more Linux users are directly connecting to the Internet. With most cable and DSL modems, you’ll need a network interface card (NIC) to connect to the modem.

When you install Fedora, it automatically detects most brands of network cards. Fedora even has the ability to detect multiple network cards in systems. The first detected network card is assigned the device name *eth0*. The second detected network card in a system (if it exists) is assigned the device name *eth1*.

Even though Fedora might have automatically detected your network card, you might still have to alter some of the configuration settings to match what your ISP requires:

1. From the GUI desktop, choose **System⇨Administration⇨Network**.
2. Enter your root password in the dialog box, if necessary.

If you're logged in as a regular user (not root), you're prompted to enter the root password.

The Network Configuration dialog box opens (refer to Figure 9-1).

3. You should see your network card listed in the Network configuration tool. Double-click the entry.

The Ethernet Device dialog box (shown in Figure 9-6) appears, showing the current configuration of the network card.



Figure 9-6:
The
Ethernet
Device
dialog box in
the Network
configura-
tion tool.

The Ethernet Device dialog box contains several settings that you can alter:

- **Nickname:** Allows you to assign a different (possibly meaningful) name to the device rather than eth0 or eth1.
- **Activate device when computer starts:** Allows Fedora to automatically activate the network card at boot up without any intervention from you.
- **Allow all users to enable and disable the device:** By default, only the root user can start and stop the network card connection. Checking this check box allows any user to do that.
- **Enable IPv6 configuration for this device:** This is a future standard for providing Internet addresses. At this time, most ISPs don't support IPv6, so leave this check box empty.

4. Select the IP address method required for your ISP.

Most cable and DSL modems provide IP addresses to devices by using DHCP. Select this method if this is your situation. If your ISP assigns a specific IP address to your device, you'll have to select the Statically set IP addresses option and enter the appropriate values.

5. Click OK to save the settings.

The Network Configuration tool saves the new network card configuration settings automatically. If the network card is not in the Active status, click the Activate button to get it started.

Look Ma, no cables!

The wireless phenomenon has taken home networking by storm. Nowadays it's just as easy to purchase and install a wireless router as it is a wired switch. If you have jumped on the wireless bandwagon, you need to know how to connect your Fedora system to your wireless access point.

Your first hurdle may be the wireless network card itself. Unfortunately, wireless network support is still in its infancy in the Linux world. Although each new version of Fedora includes support for more wireless network cards, it's still quite possible that you have a wireless network card that is not supported.



If your wireless network card is not supported, you have a last-ditch effort in the `ndiswrapper` project. The `ndiswrapper` project is an open source project that uses Windows network drivers supplied for wireless network cards to create a pseudo-Linux driver. Unfortunately, Fedora doesn't provide the `ndiswrapper` software in the default configuration, but you can find it in other areas. You can find out more about `ndiswrapper` by visiting ndiswrapper.sourceforge.net.

If you're lucky enough to have a wireless network card that is supported by Fedora, you can use our friend the Network Configuration tool to configure it for your wireless access point:

1. From the GUI desktop, choose System⇨Administration⇨Network.

2. Enter your root password in the dialog box, if necessary.

If you're logged in as a regular user (not root), you're prompted to enter the root password.

The Network Configuration dialog box opens (again, refer to Figure 9-1).

3. In the Network Configuration dialog box, click the New button.

The Select Device Type dialog box appears (refer to Figure 9-2). Because this section covers wireless Internet access, select the Wireless connection option from here, and then click Forward.

4. Select your wireless device, and click Forward.

The Select Wireless Device dialog appears, listing any detected wireless network cards on the system (shown in Figure 9-7).



Figure 9-7:
The Select
Wireless
Device
dialog box in
the Network
Configura-
tion tool.

5. Enter your wireless access point configuration settings.

The Configure Wireless Connection dialog box (shown in Figure 9-8) appears.

Figure 9-8:
The
Configure
Wireless
Connection
dialog box in
the Network
Configura-
tion tool.



You must configure several wireless settings to connect to your access point:

- **Mode:** Determines if the wireless card is managed by an access point, or is the master of other wireless cards. To connect to your access point, set this to Managed.
- **Network name (SSID):** Wireless access points are identified by a Service Set Identifier (SSID). Enter your access point's SSID in this field if you only want to connect to your wireless access point. If you leave this set to Auto, Fedora attempts to connect to any open wireless access point in range.
- **Channel:** Some wireless network cards allow you to set the specific wireless channel to use, whereas others scan all channels looking for the configured SSID.
- **Transmit rate:** Some wireless network cards allow you to specify the transmit rate, whereas others autodetect the rate used by the remote access point.
- **Key:** Encrypting wireless access points is a necessity in today's networking environment. This field allows you to enter the encryption key used by your access point.

After entering your selections, click Forward to continue.

6. Click Forward.

The IP Settings dialog box appears; refer to Figure 9-5.

7. In the IP Settings dialog box, select the appropriate radio button for whether your wireless network card needs to obtain its IP addressing information from the access point (*DHCP*, as discussed in the section “Getting Information You Need from Your ISP,” earlier in this chapter) or whether you need to tell it what its IP data is (*static*).

Most wireless access points also include a built-in DHCP server, so you can use DHCP on your wireless network even if your ISP doesn’t use DHCP.

If this is the case in your situation, select the DHCP options to automatically set all of your IP address settings. If you choose Static, the lower portion of the dialog box becomes available.

8. If you selected static, fill in the Address, Subnet Mask, and Default Gateway Address as given to you by your ISP.
9. Click Forward to proceed to the summary dialog box, and then click Apply to add the wireless network card you have just set up to your main Network Configuration dialog box.

After the wireless network card appears in the Network Configuration list, you can easily start and stop it by using the Activate and Deactivate buttons.

It’s All Fun and Games Until Something Doesn’t Work

In a perfect world, the dialup configuration steps in the preceding section would work 100 percent of the time. The Linux vendors have truly hidden all the mystery that has traditionally surrounded networking. Unfortunately, in many situations (mostly related to modems and hardware), a simplified configuration doesn’t work. If you can’t connect to the Internet after following these steps, an excellent site to find help is LinuxQuestions.org (www.linuxquestions.org). And yeah, I know — you’ll have to use a computer that already has Internet access to get that help. (Catch-22, anyone?) It’s simply impossible to anticipate the wide range of problems people can run into, and the Linux community is your best bet; this site is well known for its helpful community members.

Also, go to your favorite Web-search site and search on the error message you’re getting from the system — it doesn’t hurt to add the network hardware’s make and model and the name of your distribution as well if just the error message isn’t working. When you’re trying to figure out what’s wrong with your network connection — or trying to gather information that can help someone else figure out what’s wrong — you can use a cool tool: the handy command-line program, `ping`.



Some firewalls block the kind of traffic sent with `pings`, so these commands don't always work as expected, even with a good connection.

The `ping` command is akin to a submarine using sonar to detect other objects in the ocean. Sonar sends out a *ping* signal, which reflects off a hard surface. By measuring the amount of time between sending the ping and the ping's return, the submarine's engineer can determine whether an object is out there and how far from the submarine the object is.

The `ping` command in Linux provides information similar to what sonar provides — and it's so useful that you see it here as both a Linux command and a time-honored word. If you consider the Internet to be your ocean, you can determine, by *pinging*, what other network computers exist — and also how long it takes for your `ping` command to return. You use this command in the format `ping hostname` or `ping ipaddress`, such as `ping bob.example.com` or `ping 192.168.1.5`. If you can't ping another computer in your house or office (assuming that you have them all connected with Ethernet and not on separate dialup connections), something is wrong with the machine you're pinging from. On the other hand, if you can ping another machine in your house or office but you can't ping a machine elsewhere on the Internet, something may be wrong with your connection to your ISP.



Latency, or the amount of time it takes for a signal to travel on the Internet, has little to do with physical distance. Rather, factors such as network traffic, bandwidth, and network hardware all contribute to slow latency. These factors determine whether a `ping` command to your neighbor's computer takes longer than ping a host at the South Pole.

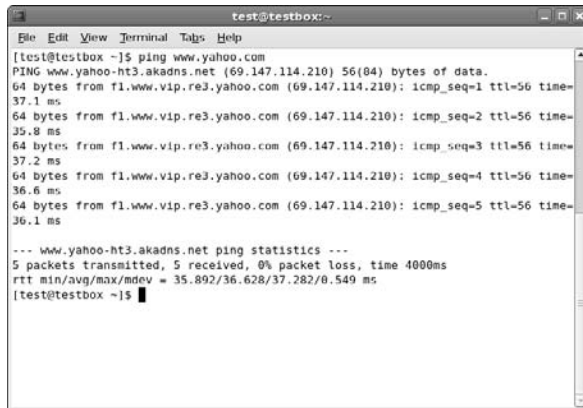
For example, try ping a Yahoo! by opening a command prompt (see Chapter 6) and typing the following command:

```
ping www.yahoo.com
```



Press `Ctrl+C` to stop the ping; otherwise, your computer continues to ping the target.

The output, shown in Figure 9-9, provides information about what `ping` is doing. If the ping can't reach the host, you receive a message that the host is unreachable. If the ping can reach the host, you receive feedback that provides how long it takes — in milliseconds (`ms`) — for the signal from your computer to get to the destination computer and back again. (The lower the numbers, the better.) For computers connected through Ethernet, a ping time of 1 ms to 3 ms is an acceptable response time. For dialup connections, expect somewhere around 150 ms. When you start seeing ping times climbing to 900 ms or higher, the network is likely under heavy use (or you have something wrong with your cables).



```
test@testbox:~  
File Edit View Terminal Tabs Help  
[test@testbox ~]$ ping www.yahoo.com  
PING www.yahoo-ht3.akadns.net (69.147.114.210) 56(84) bytes of data:  
64 bytes from f1.www.vip.re3.yahoo.com (69.147.114.210): icmp_seq=1 ttl=56 time=37.1 ms  
64 bytes from f1.www.vip.re3.yahoo.com (69.147.114.210): icmp_seq=2 ttl=56 time=35.8 ms  
64 bytes from f1.www.vip.re3.yahoo.com (69.147.114.210): icmp_seq=3 ttl=56 time=37.2 ms  
64 bytes from f1.www.vip.re3.yahoo.com (69.147.114.210): icmp_seq=4 ttl=56 time=36.6 ms  
64 bytes from f1.www.vip.re3.yahoo.com (69.147.114.210): icmp_seq=5 ttl=56 time=36.1 ms  
--- www.yahoo-ht3.akadns.net ping statistics ---  
5 packets transmitted, 5 received, 0% packet loss, time 4000ms  
rtt min/avg/max/mdev = 35.892/36.628/37.282/0.549 ms  
[test@testbox ~]$
```

Figure 9-9:
Sample
results of
the ping
command.

After You're Connected

After you're connected, you'll probably want to leap straight to Chapter 10 for how to use the various Internet programs available to you. In the meantime, let me offer you a number of sites you might find useful when you need help, have questions, or otherwise want to explore the Linux world farther — some of these are pretty geeky, but you'll grow into them! You also may want to update your machine (see Chapter 16), add new software (also Chapter 16), and give extra thought to system security (see Chapter 18) to protect yourself from nasty people and programs.

Useful sites include

- ✓ www.linuxsecurity.com: A somewhat geeky (but useful) place where you can track what's going on in the world of Linux security.
- ✓ www.kerneltraffic.org: Everything you never wanted to know about what's happening with the Linux kernel but were afraid to ask. This place is great for tracking what new features are being added, if you're interested in things like bleeding-edge hardware or just like to be "in the know."
- ✓ www.linuxquestions.org: I mention this site elsewhere in the book, but it's worth mentioning again. It's a popular hangout for people who have questions to ask, and for those who like to help them find answers.
- ✓ <http://linux-newbie.sunsite.dk/>: The Linux Newbie Administrator Guide is a massive collection of information that you may find useful.

- ✓ www.tuxfiles.org: A collection of short help files on individual topics of interest.
- ✓ www.tldp.org: Contains lots of help documentation written for various user levels.
- ✓ www.tuxmagazine.com: A free digital magazine for Linux desktop users.
- ✓ www.slashdot.org: Uber-geek hangout (with a rather hostile culture, unfortunately) that features lots of pointers to interesting articles online. Best used to view the articles; skip the arguments in the comments section.

In addition to these recommendations, keep in mind that your favorite search site can also be very handy.

Chapter 10

Using the Internet

In This Chapter

- ▶ Browsing the Web
 - ▶ Checking and sending e-mail
 - ▶ Grabbing files
 - ▶ Sending instant messages
 - ▶ Talking across the Internet
-

Give a man a fish and you feed him for a day; teach him to use the Net, and he won't bother you for weeks.

— Anonymous

The *Internet* is a vast network of computers that spans the globe. Many different types of computers and operating systems work together to allow you access to information across the Internet. Linux, along with the other related UNIX operating systems, has long supported and worked with the Internet. Practically all the different services available on the Internet are available from your Linux desktop.

When you install Linux (see Chapter 3), one or more Web browsers, mail programs, and instant messaging tools are placed onto your new system. In this chapter, I introduce you to some tools you can use to access different services on the Internet, such as accessing Web sites, using e-mail, investigating newsgroups, and utilizing FTP — assuming that your Internet connection is configured (see Chapter 9).

Browse the Web with Firefox

Many people attribute the explosive growth of the Internet to the graphical Web browser. The Internet has been around for much longer than the invention of the browser. It's just that most of the work done on the Internet was in

plain old text, which held little attraction for those people who like pretty pictures. In the world of Linux, the most popular browser is arguably Firefox (www.mozilla.org/products/firefox), which is based on Mozilla — essentially Firefox is just the Web browser portion of Mozilla, whereas Mozilla can also handle e-mail and news browsing. Another popular option is Konqueror (www.konqueror.org), which besides being a good Web browser, also includes file-browsing features similar to the Nautilus program.

Configuring Firefox

You can start surfing right now, if you want. You don't need to customize your browser. However, you may want to take a moment to tell Firefox your preferences, such as the default Web site to show when it starts up, what font sizes to use by default, what colors to use, and many other options.

The following steps introduce you to the Preferences window, where all the Firefox configuration parameters are stored:

1. Start Firefox.

Fedora provides two ways for you to start Firefox. You can use the standard menu entry by clicking Applications ⇨ Internet ⇨ Firefox, or, if you're in a hurry, you can just click the globe icon in the top Panel.

2. Choose Edit⇨Preferences.

The Preferences window appears, as shown in Figure 10-1. This dialog box contains all the settings you need to play with for your Web browsing experience.

3. Click the icons along the top of the Preferences window to access the various preferences categories.

The Advanced category is complex enough to offer tabs for various sub-categories. These tabs can be worth exploring as well.

4. If you get tired of reading through menus, just click Close to close the Preferences dialog box and get back to surfing.

General preferences

Selecting the Main option icon — shown earlier in Figure 10-1 — allows you to determine how you want Firefox to look and act when it runs. Two separate sections govern this dialog box:

- ✓ **Startup:** This section allows you to choose what Firefox shows you immediately upon starting. You can type a URL directly into the Home Page text box, or click one of the following buttons:




Figure 10-1:
The Firefox
Preferences
window in
Fedora 7.

- **Use Current Page:** Navigate to the page you want to use for your default Web page, and then click this button. The URL for this page appears in the Location(s) text box.
- **Use Bookmark:** Click this button to open the Set Home Page dialog box, which contains all your bookmarks. Select the bookmark you want to use and then click OK to add the URL for this page into the Location(s) text box.
- **Restore to Default:** Click this button to have Firefox open to just an empty window.

✓ **Downloads:** This section allows you to set how Firefox performs file downloads. You can set where you want Firefox to download files to, as well as set if the Downloads window appears while a file is downloading.

Privacy preferences

Every time you send an e-mail or click a Web site, you're exchanging information across a network. Although you don't need to lose sleep over it, you should be conscious of every request you make on the Internet. Take a moment to examine the Privacy preferences (see Figure 10-2) and determine if there's anything you want to change:

Figure 10-2:
The Firefox
Privacy
Preferences
window in
Fedora 7.



- ✓ **History:** The History section lets you designate how many days you want your browser to remember where you've been. You can also choose to let Firefox remember text you've entered into forms and the search bar. If you use this feature, Firefox will attempt to autocomplete words you've already typed before in those forms. You can also allow Firefox to remember the files you've downloaded so you can easily review them instead of accidentally downloading a 500MB file that you already downloaded last week.
- ✓ **Cookies:** This is one of those buzzwords that sends some people into a frenzy. Cookies are bits of information a Web site stores on your machine so that this data can be accessed again when you return. Some people don't care how many cookies are created and some don't want any on their systems. To prevent cookies from being stored, click the Accept Cookies from Sites box to uncheck it. You can also only allow certain sites to set cookies by clicking the Exceptions button and adding the sites into the list.

You can view every cookie you currently have by clicking the Show Cookies button, and can either remove them individually (select the cookie you want to remove and then click Remove Cookie) or click the Remove All Cookies button to toss out all of your cookies immediately. For the Keep Until drop-down list box, choose how often you want Firefox to dump cookies:

 - **Until they expire:** Most cookies have an automatic expiry date. If you want to let the cookies dictate when they will be deleted, choose this option.

- **I close Firefox:** If you are using a library computer or another machine that is shared, or just don't want the computer storing cookies past when you're done using the browser, choose this option.
 - **Ask me every time:** If you want to be ultra-paranoid about cookies or to find how often they are actually used, choose this option.
- ✓ **Private Data:** This section allows you to set what stored items (such as cookies, temporary files, history items, and stored form data) are cleared, and how they are cleared. Select the Always Clear My Private Data When I Close Firefox check box if you want to clear all of your stored files every time you exit Firefox. You can select exactly what items are cleared by clicking the Settings button and choosing the individual items that you want cleared. As a safety feature, the Ask Me Before Clearing Private Data check box is selected by default. This helps prevent accidentally clearing your data.

Content preferences

Selecting Content brings you to the Content portion of the Firefox Preferences dialog box (Figure 10-3). This section lets you control how Firefox reacts to various types of content. The top portion refers to items to allow through or not — notice that each of these has a button to its right that offers additional control:



Figure 10-3:
The Firefox
Content
Preferences
window in
Fedora 7.

- ✓ **Block Popup Windows:** Tells Firefox to block new windows that you didn't ask for — this is a favorite feature and one reason that many people move to Firefox under Microsoft Windows. Some sites have functionality that require the ability to open, say, login windows, so you can click the Allowed Sites button to make exceptions for them. If you don't configure this feature, it still works. A yellow bar appears in the browser window whenever a pop-up is blocked, allowing you to access the settings at that point.
- ✓ **Load Images Automatically:** Lets you block all images. Useful for those who are sight impaired or who are trying to speed up browsing by ignoring pictures. Selecting For the Originating Web Site Only blocks many ads because ads are often loaded from other sites and just appear on the one you're viewing, but it can also block legitimate graphics. The Exceptions button helps you manage what to see and what not to.
- ✓ **Enable Java:** If you have Java installed and this option enabled, Firefox runs Java code it finds.
- ✓ **Enable JavaScript:** Some people shut off this feature for security reasons. Many don't. Clicking the Advanced button opens the Advanced JavaScript Settings dialog box, which you can use to choose what you will allow this scripting language to do. If you find a site that does annoying things like resize your windows, go to this dialog box to disable that ability.

The middle section is Fonts & Colors. These features include

- ✓ **Default Font:** Click this drop-down list box to select the font you want to use.
- ✓ **Size:** Click this drop-down list box to choose the font size you want to use.
- ✓ **Advanced:** Click this button to open the Fonts dialog box, which lets you choose the default fonts and sizes to use for various font classes, along with other visual settings. Here you can also set that pages can use their own fonts rather than your overrides.
- ✓ **Colors:** Click this button to open the Colors dialog box and change the default colors assigned to text, the background, and links.

The lower section is File Types. Clicking the Manage button allows you to assign specific external programs to handle specific types of files (for example, you can configure OpenOffice.org to handle .doc files).



When browsing, if you are unhappy with a font's size, you can hold down the Ctrl+ keys to make the fonts bigger and the Ctrl– (minus) keys to make the fonts smaller. The + and – are the versions on your number keypad, not on the main keyboard.

Advanced preferences

This section provides four tabbed sections of settings for more advanced features that you probably won't need to mess with:

- ✓ **General:** The General tab lets you set up features for the movement impaired, extra browser control features, and the languages you can read in the order of preference.
- ✓ **Network:** Click the Settings button in the Connections section if your network administrator told you that you have to use *proxy servers* to access online content. The Cache section allows you to set how much disk space Firefox is allowed to consume when storing visited Web page content. This allows Firefox to display the pages faster the next time you visit them. You can clear the cache by clicking the Clear Now button. The cache data is also part of the private data that you can set to clear automatically when you exit Firefox.
- ✓ **Update:** The Update tab tells Firefox how to handle checking for updates, both for itself, and for any plug-ins you have installed.
- ✓ **Encryption:** The encryption tab has settings for security protocols Firefox uses, and if you have any special encryption certificates for visiting encrypted Web sites. Usually, you can leave these settings alone.

Them dad-gum browser plug-ins

Do you ever get irritated at those Web sites that insist that you download a *plug-in*, or additional piece of software, just to view the site? The difference between a plug-in and an external program is this: A plug-in displays the results in the browser, and an external program runs outside the browser. Although these plug-ins are annoying if you're just looking for some basic information, they can provide some pretty cool stuff, such as streaming video and music through your Web browser.

Chapter 16 explains how to add the popular plug-ins (like Macromedia's Flash plug-in) to your system. When you download a particular type of file, Firefox gives you the chance to tell it how to automatically handle those types of files later. Beware of telling the browser to run things automatically: You could be opening yourself up for security problems.

Firefox also supports a concept called *extensions*, which refers to extra features that aren't included in the main browser because many people may not want to bother with them. Select Tools⇨Add-ons to open the Add-ons dialog box (see Figure 10-4).

Figure 10-4:
The Firefox
Add-ons
dialog box in
Fedora 7.



The Extensions section shows the currently installed extensions. Click the Get More Extensions link (if you are connected to the Internet) to open a Web page that lets you surf what extensions are available. If you find one you want to try, click the Install link underneath its description and then click Install Now if your version of Firefox falls in the listed range of versions supported (choose Help→About Mozilla Firefox to find out). Some extensions may give you an error and may be uninstallable because of it, unfortunately.

Surfing the Web

Firefox is your viewing window into the wonderful World Wide Web. Firefox's primary purpose is to fetch Web pages on your command, download all their graphics and related files into your computer's memory, and, finally, render the page for your interactive viewing pleasure.

If you're used to using Internet Explorer or Netscape, using Firefox should be a snap. It has all the familiar navigation tools, such as an address bar; Back, Forward, Reload, and Stop buttons; and a feature that stores links to your favorite Web sites (bookmarks).

Firefox, like Netscape, also has a History sidebar feature — press Ctrl+H to show or hide it (see Figure 10-5) — which you can use to access a Web page you recently visited.

Tabbed browsing — a feature of Firefox that may be new to you — enables you to open several Web sites in the same window. For many people coming from Internet Explorer 6 or Netscape, the tabbed browsing feature of Firefox may seem odd at first, but I think you'll quickly fall in love with it.




Figure 10-5:
Firefox with
the History
Sidebar
open, in
Fedora 7.

After you start the Firefox Web browser, the Web site defined as your home page appears in a tabbed page. You can keep that page open and open a new Web site in a new tabbed page by pressing Ctrl+T. The new Web page appears within the same Firefox window, but as a tabbed page. You can switch between the two tabbed pages by simply clicking the appropriate tab at the top of each page.

You can also open Web links as new tabbed pages by right-clicking the link and selecting Open Link in New Tab. This feature is great for navigating through complex Web sites where you need to refer back to previous pages.

Evolving into E-Mail

Although most people think of the Web when they think of the Internet, e-mail may, in fact, be the most used and beloved of Internet applications. A multitude of e-mail programs are available for Linux users. Fedora chose to include the Evolution e-mail package by default.

Evolution should remind you very much of Outlook. If you like being able to integrate your calendar, address book, task manager, and e-mail, you should feel right at home with Evolution. Other Linux e-mail clients such as Mozilla Mail, Thunderbird, and KMail are more single-purpose, although Thunderbird also supports news reading. If you like to keep things simple, you might want to try one of these packages.

Setting up Evolution

The first time you open Evolution, the program walks you through the setup process, so be sure to have the following information on hand:

- ✓ Your assigned e-mail address, in the format `username@example.org`.
- ✓ The type of mail server used for incoming mail: for example, POP.
- ✓ The name of the mail server used for incoming mail, in the format `servername.example.org`.
- ✓ The type of mail server used for outgoing mail: for example, SMTP.
- ✓ The name of the mail server used for outgoing mail, in the format `servername.example.org`. This name may be the same as for the incoming mail server.
- ✓ Any special type of *authentication* required, for both mail coming in and going out. This authentication is typically just passwords for incoming mail and may be nothing for outgoing mail.

After you start Evolution for the first time, the Evolution Setup Assistant launches.

Click Forward to proceed past the Welcome screen, after which the Identity dialog box (see Figure 10-6) opens. To fill in this dialog box, complete the following steps:



The screenshot shows a window titled "Evolution Setup Assistant" with a tab labeled "Identity". Inside the window, there is a text area with the instruction: "Please enter your name and email address below. The *optional* fields below do not need to be filled in, unless you wish to include this information in email you send." Below this, there are two sections: "Required Information" and "Optional Information". Under "Required Information", there are two text boxes: "Full Name:" with the value "Ima Test" and "Email Address:" with the value "itest@example.org". Under "Optional Information", there is a checkbox labeled "Make this my default account" which is checked. Below the checkbox are two more text boxes: "Reply-To:" and "Organization:", both of which are empty. At the bottom of the window, there are three buttons: "Cancel", "Back", and "Forward".

Figure 10-6:
The
Evolution
Identity
dialog box in
Fedora 7.

1. Change the Full Name field, if you want.

Some people don't want to use their real name and opt to use a nickname. There may or may not be default values assigned, depending on what other programs you've configured on your machine.

2. Change the E-Mail Address field if it doesn't match the address your ISP gave you (it probably doesn't).

3. If you want this particular account to be your default mail account, make sure the Make This My Default Account check box is selected.



If you want people to respond to a different e-mail account than what you're sending the message from, enter that e-mail address in the Reply-To field. Otherwise, you can keep this field blank. You can also choose to enter your organization in the Organization field.

4. Click Forward to proceed.

The Receiving E-Mail dialog box opens.

5. In the Server Type list box, select the type of incoming e-mail server your ISP uses.

Leave it as None if you don't want to receive e-mail on this computer — this setting is useful if you only want to be able to send or only want to use the calendar and task management features in Evolution. Depending on which item you choose, the dialog box changes to ask for the appropriate information. I assume that you're using a POP mail server to receive mail, so what you'll see when you select POP is shown in Figure 10-7. If you're not using POP, complete the dialog box you do see and then proceed to Step 17.

Evolution Setup Assistant

Receiving Email

Please select among the following options

Server Type: POP

Description: For connecting to and downloading mail from POP servers.

Configuration

Server: servername.example.org

Username: itest

Security

Use Secure Connection: No encryption

Authentication Type

Password

☒ Remember password

Figure 10-7:
The
Evolution
Receiving
E-Mail
dialog box,
with POP
selected, in
Fedora 7.



The latest versions of Evolution have support for Microsoft Exchange servers available. If you need this, it may or may not be already installed in your distribution. Use the information in Chapter 16 to add the “evolution-connector” package to your system and then restart the Evolution setup process.

6. Enter the full name for your POP mail server in the Server text box.

The name may be something like `pop.example.com`.

7. Enter your login name for checking mail in the Username text box.

If your e-mail address is `jane@example.com`, your username is `jane`.

8. If you were told to use SSL for security, in the Use Secure Connection (SSL) drop-down list box, select either Always or Whenever Possible, depending on what your ISP has specified.

9. Under Authentication Type, select the appropriate authentication option.

If you don't know what kinds of authentication your mail server uses, click the Check for Supported Types button. If your ISP server supports an encrypted authentication method, Evolution will automatically detect and use it. Typically, the authentication method is just a plain text password, so you can select the Password option.

10. Check the Remember This Password check box so that you don't have to enter your e-mail password every time you check mail.

If you're going to walk away from your computer and someone else may possibly access it, protect yourself from outgoing prank e-mails by going to the Main Menu and selecting Lock Screen. Doing so makes sure that no one can use your GUI until you type in your login password.



You have to have your screen saver turned on for this feature to work. Screen savers are on by default.

11. Click Forward to proceed to the secondary Receiving Options dialog box (see Figure 10-8).


Whether this dialog box exists is determined by which type of incoming mail server you're using. If you're not using POP and you get a secondary dialog box, it may not match the one shown here, although it does share some of the options shown.

12. If you want to check for new e-mail automatically, check the Automatically Check for New Mail check box.

If you don't have a permanent connection to the Internet, you may prefer to check mail manually. If so, leave the box unchecked and skip to Step 14.

13. In the Minutes text box, set how often you want to automatically check for new e-mail.

Figure 10-8:
The
Evolution
Receiving
Options
dialog box,
for POP mail
users in
Fedora 7.



14. **If you check mail from multiple computers, you may want to check the Leave Messages on Server check box so that you can access the same messages from all your machines.**

If you select this option, you'll end up with copies of messages on multiple machines, but it also means that you'll have access to your e-mail no matter where you're checking it from. Keep in mind that occasionally you'll need to uncheck this box and check mail, just to clear out all the space that your e-mail is taking up on the server. You don't want to run out of space.

15. **If your ISP told you not to use POP3 extensions, check the Disable Support for All POP3 Extensions check box.**

16. **Click Forward to proceed.**

You're now finished setting up for incoming mail. In the Sending E-mail dialog box, shown in Figure 10-9, do the following:

17. **Change the Server Type entry if yours isn't SMTP.**

I assume that you're sending e-mail with SMTP.

18. **Enter the full name of the SMTP mail server in the Host text box.**

The name may be something like `smtp.example.com`.

19. **If you were told to use SSL for security, in the Use Secure Connection (SSL) drop-down list box, select either Always or Whenever Possible, depending on what your ISP has specified.**

Figure 10-9:
The
Evolution
Sending
E-mail
dialog box in
Fedora 7.

- 20. If you were told to use additional authentication for sending mail, select the Server Requires Authentication check box.**

Typically, nothing is required here. If you don't check this box, skip to the next set of instructions. Otherwise, proceed to the next step.

- 21. Under Authentication, select the appropriate Type as instructed by your ISP.**

If you're not sure, click Check for Supported Types.

- 22. Under Username, enter the login name you're supposed to use to authenticate with your outgoing mail server.**

This name may be different from the one you use for incoming mail.

- 23. Select the Remember Password check box so that you don't have to enter your password each time you send mail.**

Just make sure, again, to log out or use the Lock Screen option so that no one can send joke e-mails at your expense!

- 24. Click Forward to proceed.**

You have now reached the wonderfully simple Account Management dialog box.

- 25. Either leave the Name entry as it is (your e-mail address) or change it to something descriptive so that you can tell which account you're looking at if you have a list of accounts you need to use.**

It's common for your e-mail address to appear in this final box as something like `jane@computer5.example.com`. If this happens for you,

you'll want to edit this entry so that it matches your real e-mail address, which in Jane's case is `jane@example.com`.

26. Click Forward to proceed to the Time Zone dialog box.

Click the city closest to yours on the time zone map, or use the Select drop-down list bar to choose your proper time zone.

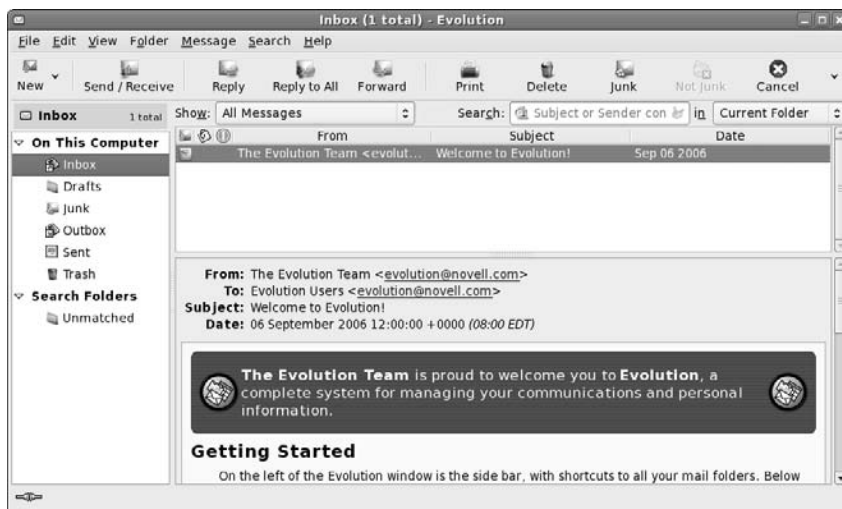
27. Click Forward to proceed to the Done dialog box.

Here you see just a quick message saying that you're finished.

28. Click Apply to save your settings.

The Evolution program opens, as shown in Figure 10-10.

Figure 10-10:
The
Evolution
application
window in
Fedora 7.



Sending and checking e-mail

The following steps outline how to create a new e-mail message and send it:

1. Click New.

A Compose a Message window opens, as shown in Figure 10-11. If you want to open another new item such as a new contact list entry or a new calendar appointment, click the arrow next to New to open the list of options to choose from.

2. Type the recipient's e-mail address (such as `bob@example.net`) or a list of addresses separated by commas (such as `bob@example.net, tom@example.org`) in the To text box.




Figure 10-11:
An Evolution
Compose
Message
window in
Fedora 7.

3. If you need to add a CC (Carbon Copy) or BCC (Blind Carbon Copy) to the list of recipients, open the mail message's View menu and select the appropriate field(s) to appear and then enter the appropriate address(es) into those fields.

Don't be confused on the CC and BCC options. Mail addresses added to the CC field appear in the header of the mail message for everyone to see. Mail addresses added to the BCC field don't appear in the header of the mail message (but the message is still delivered to them). This is a sneaky way to send someone a copy of a message without others knowing it.

4. Type the topic of your e-mail into the Subject text box.
5. In the lower window, type the body of your e-mail.

Use the handy formatting buttons and the Format menu if you want to "pretty up" your e-mail. If you want to use the formatting buttons, set your Format type to HTML first.

6. If you want to add a signature to the bottom of your e-mail, click the drop-down list box next to Signature and select Autogenerated.

To create custom signatures, go to the main Evolution window and choose Edit→Preferences to open the Evolution Preferences dialog box (see Figure 10-12). Select the account you want to create the signature(s) for and then click the Edit button to open the Evolution Account Editor. In this dialog box, in the Identity tab, click Add New Signature. Here, you can create and format your new signature, and after you click the Save and Close button, you can choose which signature should be the default in the Identity tab's Signature drop-down list box. Get rid of the extra dialog boxes by clicking OK or Close on each one.

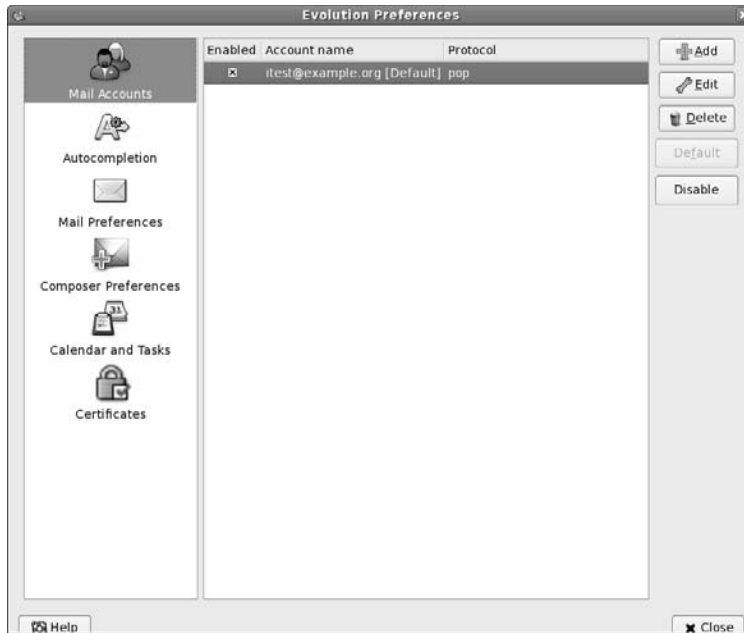


Figure 10-12:
The
Evolution
Preferences
dialog box in
Fedora 7.

7. When you finish typing your message, click the Send button.

The e-mail is now added to your Outbox, where you can edit it if you want.

8. Click the Send/Receive button.

Your e-mail goes out, and Evolution checks for new incoming mail.

Take some time to really explore Evolution. As you can see just from the figures in this chapter, this program has many features, including the ability to filter junk mail.

Taking Advantage of Instant Messaging

Instant messaging (IM) between people is like using a telephone — except that you type your conversation rather than speak it. In addition, you can simultaneously hold multiple instant messaging conversations without the need for additional connections to the Internet. America Online (AOL) provides one popular instant messaging service, named AOL Instant Messenger, or AIM. Others are ICQ, MSN, Yahoo!, and Google. A wide variety of computer operating systems, including Linux, support these various services. Fedora 7 uses the Pidgin instant messenger client, which allows you to connect to any of these instant messaging services.




Until recently, the Pidgin software package was called GAIM. Unfortunately, the name of the package had to be changed due to copyright restrictions. If you've used GAIM in previous versions of Fedora, Pidgin works exactly the same way.

Using the Pidgin Instant Messenger

To launch Pidgin, choose Applications⇨Internet⇨Internet Messenger. After the application starts for the first time (see Figure 10-13), its main dialog box (called the Buddy List) and the account management dialog box both appear. To tell Pidgin about one of your IM accounts:

Figure 10-13:
The Pidgin
IM client
with its
Buddy List
plus
Accounts
dialog box
open in
Fedora 7.



- 1. Click the Add button in the Accounts dialog box.**

The Add Account dialog box, shown in Figure 10-13, appears.

- 2. In the Protocol drop-down list box, select the IM network you want to use.**

Your options are AIM, Gadu-Gadu, GroupWise, ICQ, IRC, Jabber, MSN, QQ, SILC, SIMPLE, Yahoo, and Zephyr.

If you want to use Google Instant Messenger, see the instructions at www.google.com/support/talk/bin/answer.py?answer=24073.

- 3. In the Screen Name text box, enter the login name for your IM account.**

If you're using ICQ, use your membership number.



4. In the Password text box, enter your IM password for this account.
5. In the Alias text box, enter the name you actually want to show in people's IM clients, unless you want to use your screen name.
6. Select the Remember Password check box to put a check in it so that you don't have to enter your password every time you connect to this IM service.

Remember to log out or use the Lock Screen option if you walk away from your computer, just in case someone gets it in his head to go play a trick on you and send messages to people by using your IM client.

7. If you get e-mail through this service and want to know when new mail has arrived, check the New Mail Notifications check box.
8. If you want to use a tiny picture as a *buddy icon* (of yourself, your dog, or whatever else you might want to use), click the Open button next to the Buddy Icon label and then navigate to the picture you want to use.
9. If you want access to the more advanced options for this IM service, click the Advanced tab.

Each IM service has its own set of advanced options you can configure. For example, for the AIM IM service you can change the default AOL server to connect with and the default TCP port number.

10. When you finish entering your information, click Save to add this IM account to your accounts list.

Go through this process for each account you want to use with Pidgin. Then in the main dialog box, click the box in the Enabled column for the accounts you want to automatically log in.

You can modify the account settings at any time by selecting the account in the main dialog box and then clicking the Modify button.

After you start a session, the Buddy List window displays your active connections. If you close the Buddy List window, the Pidgin icon appears on your Panel as the program runs in background. You can open the Buddy List window by right-clicking this icon.

Again, Pidgin has more to it than this, so it's definitely worth playing around with this tool.



After you close the Accounts and Buddy List dialog boxes, you can reopen them by right-clicking the AIM icon in your system tray and selecting the dialog box you want to open from the menu.

Troubleshooting your IM connections

If you've been using most of these IM clients with other operating systems, after your Linux system has successfully connected to the IM service, your existing buddy lists are automatically imported. If you haven't ever used AIM or MSN, you can now set up buddy lists. *Buddy lists* contain the usernames of people you want to communicate with through the instant messaging service. Your buddy list lets you know when your "buddies" are online and available to receive an instant message.

Sometimes, your Linux system is unable to connect properly to the IM service. When this problem happens, an error message pops up on your screen and indicates a failure to connect. You may be unable to connect to the IM service for several reasons:

- ✔ You may have entered the wrong password for your IM account or chosen the incorrect IM account name.
- ✔ Your computer may not be connected to the Internet. Try opening your Web browser to see whether you can get to a Web site, which tells you whether you're connected to the Internet.
- ✔ If you can open a Web site but can't get IM to work, the IM system may be unavailable. This problem occurs at times because of maintenance of the IM service or an excessive amount of traffic on the Internet or on the IM service.
- ✔ If you attempt to access the IM service from your computer at work, your company or organization may block the IM service for security or productivity reasons. If using IM at your work is permitted, check with your network administrator to see whether he can help you out.
- ✔ Often, companies use firewalls between the company's network and the Internet to keep out unwanted traffic on the company's network. If the firewall is configured to block IM traffic, you cannot use IM across the Internet.

Downloading with BitTorrent

Say what you want about file sharing, but it's popular. Networks such as BitTorrent (www.bittorrent.com) have actually found legitimate use among software companies and other content distributors, as a way of offering downloads for larger files without having to take the brunt of the bandwidth use themselves, so it's worth looking at how to use this software under Linux.


The standard Fedora installation doesn't include a BitTorrent client by default, so you'll have to install it yourself if you want to use it. You can use the information in Chapter 16 to install the `bittorrent-gui` package.

After installing the BitTorrent client you can start it by choosing Applications⇨Internet⇨BitTorrent File Transfer. When you find out that a file is available over BitTorrent, you first download the `.torrent` file for the item. After you have this file downloaded:

1. **Start your BitTorrent client as discussed earlier.**

The Fedora client is shown in Figure 10-14.

Figure 10-14:
The
BitTorrent
window in
Fedora 7.



2. **Select File⇨Open Torrent File.**


A file browser window appears.

3. **Navigate to and select the `.torrent` file for the document you want to download, and then click OK.**

The file appears in what looks a lot like a music player format, as shown in Figure 10-15. You can speed up the download by offering faster uploads to other BitTorrent users, using the Maximum Upload Rate slider. If you have broadband (see Chapter 10), 40 kbps works well.

It can take a minute or so for the tool to fully synchronize, so expect the projected time to drastically change for a bit until it settles on a consistent value.

Figure 10-15:
The
BitTorrent
window in
Fedora 7
with a file
down-
loading.



4. After the download finishes, leave BitTorrent running to allow others to retrieve parts of the download file from you. When you're ready to exit, choose File→Quit to close the client.

BitTorrent creates a folder on your desktop that contains the downloaded files. Even after the download finishes, BitTorrent continues to host the file for others to download. If you want to be a good Internet neighbor, keep the BitTorrent application running for a while after your download finishes to let others grab their pieces as well.

Ol' Fashion File Transfers

While BitTorrent is the wave of the future, there are still plenty of Internet sites that require you to download files by using the old File Transfer Protocol (FTP) method. You will often find file repositories for applications, utilities, and other neat stuff on FTP sites. Fedora includes a command line ftp client (simply called ftp), but doesn't include an FTP client by default, but you can easily add one.

The gFTP program is a popular FTP client for the Linux environment. You can use the standard techniques discussed in Chapter 16 to download and install the gFTP package. After it's installed, you can go exploring:

1. Click Applications→Internet→gFTP.

The main gFTP window appears, as shown in Figure 10-16. By default, gFTP points to your home directory as the local directory (shown in the left-side list box).




Figure 10-16:
The gFTP
main
window.

2. Enter the hostname of the remote FTP server in the host drop-down box, the username in the user drop-down box, and the password in the pass drop-down box.

If you've connected to a host previously, its hostname and user login name appear in the drop-down boxes. You can select these names from the list.

3. When the remote FTP session starts, a list of files and folders available appears in the right-side list box.
4. Find and select the file(s) you want to download, select the folder to download them to in the left-side list box, and then click the left-pointing arrow to transfer the file(s).
5. When finished, click the icon next to the hostname to disconnect and end the FTP session.

It's as simple as that. Using a graphical FTP client makes transferring lots of files almost painless.

Talkin' on the Phone

The new craze these days is chatting on the phone over the Internet. Whether you're just wanting to experiment or are a complete devotee, Linux offers the software you need to take part. Think of the Internet phone networks like Instant Messenger networks. One of the most popular networks for supporting Internet phones is Skype (www.skype.com), which uses SIP (Session Initiation Protocol), a popular default that you can use a wider variety of programs with. You first need the proper hardware, whether it's a phone you can actually plug into your computer (a SIP phone, for example) or just a combination of a headset and a microphone — preferably not the lowest-end model so you'll get the best sound quality. When you have the hardware, it's time to get the software.

In Fedora, the default client is Ekiga (www.ekiga.org). It is available through Applications→Internet→IP Telephony. Two other SIP clients are also included in the repositories, as discussed in Chapter 16. They are KPhone (<http://sourceforge.net/projects/kphone>) and Linphone (www.linphone.org). The popular Skype package isn't included, but you can download a Linux client from www.skype.com.

In the upcoming sections, I show you how to configure and use both the default Ekiga package and the popular Skype package.

Using the default: Ekiga

The default Ekiga Internet Phone software is a redesign of the old GNOME GnomeMeeting package for Internet voice and video conferencing applications. Ekiga was expanded to provide the same basic functions as the more popular Skype software, along with compatibility with Microsoft NetMeeting. This compatibility allows you to join in conferences with your Microsoft friends (just don't brag too much that your software was free).

To start Ekiga, click on Applications→Internet→IP Telephony, VoIP and Video Conferencing. The first time you start Ekiga, you get a series of configuration windows. Follow these steps to get your Ekiga software configured:

1. Click Forward to start the wizard.

The Personal Information dialog box appears.

2. Where prompted, type the name you want to be identified by on the network.

Other Ekiga users see you listed by the name you enter here.

3. Click Forward.

The wizard proceeds to the Ekiga Account page.

4. Type your Ekiga account username and password where prompted.

If you don't have an account yet, you can jump over to the Ekiga Web site (www.ekiga.net) and register for a free account.

5. Click Forward to continue.

6. In the resulting Connection Type window, select the type of Internet connection your system uses and then click Forward.

The configuration wizard is now set to go through a series of windows to automatically determine hardware and software specifics of your system setup.

7. For each subsequent window, click Forward after the relevant determination or detection is completed.

The series of windows includes the following:

- **NAT Type:** Determines whether a Network Address Translation (NAT) server is in place between your system and the Internet. Remote Ekiga users wanting to place a call to you must be able to access your system from the Internet. If you use a NAT for Internet connectivity, your network address is the address of the NAT, not your system.

- **Audio Manager:** Detects the software that controls your audio applications. For most Linux Distributions, this software is called ALSA.
- **Audio Devices:** Detect the audio input and output devices on your system. The Ekiga software will most likely succeed in detecting these devices with no intervention on your part.
- **Video Manager:** Detects the video manager software on your system. For most Linux systems, this software is Video4Linux (V4L).
- **Video Devices:** Determines whether a video input device (webcam) is connected to your system.

When the configuration wizard is complete, a final configuration window appears, showing the choices you've made. Click the Apply button to finalize the configuration.

After finishing the configuration, the main Ekiga window (shown in Figure 10-17) appears, listening for new calls and waiting for you to place a call.



Figure 10-17:
The main
Ekiga
window.

Talking with Skype

Skype is by far the most popular IP phone network around. Many Linux distributions install Skype as the default IP phone software. Unfortunately, Fedora doesn't. If you prefer to use Skype on your Fedora system, you have to install it yourself.

To install Skype in Fedora, do the following:

1. **Point your Web browser to `www.skype.com`.**

2. **On the Skype home page, click the Download link.**

You are taken to the Download Skype page.

3. **Click the Skype for Linux link on the left side of the page.**

You are taken to the Download Skype for Linux page.

4. **Right-click the Download the Public Key link and select Save Link As from the context menu that appears.**

The Save As dialog box appears.

5. **Click Save to save the file into your download directory (by default, this is your desktop).**

The dialog box closes, and the file appears on your desktop.

6. **Select Applications⇨System⇨Terminal.**

A terminal window appears, as discussed in Chapter 6.

7. **Type `su -` to temporarily become the root user.**

8. **Type `rpm --import /home/user/Desktop/rpm-public-key.asc`, where *user* is your user account, such as `/home/dee/Desktop/rpm-public-key.asc`.**

I am assuming that your downloads directory is desktop. If it isn't, adjust the file path accordingly.

9. **Click the X in the upper right corner of the terminal window to close it.**

10. **Back on the Skype Download Web page, click the RPM for Fedora entry, even if it isn't for Fedora 7.**

Notice that you can get Skype for other Linux distributions here too. An Opening Skype dialog box appears, with the full name of the file in the title.

11. **Leave Open with Software Installer selected and click OK.**

The Downloads window opens during the download, and then when the download is complete, the Password For root dialog box opens.

12. **Enter your root (administrator) password and click OK.**

The Installing Packages dialog box appears.

13. **Click Apply.**

The software installer adds the software to your system. When the installation is finished, an Information dialog box appears.

14. Click OK to close the dialog box, and click the X in the upper right corner of the Downloads window.
15. If Skype doesn't appear in your Applications⇨Internet menu, log out and back in (no need to reboot). It should show up after that.

When Skype starts, a wizard appears, allowing you to either enter your Skype account information or connect to a Skype server and create a new account. After you have an account, you can use the Skype interface, as shown in Figure 10-18. Before being able to make a call, you need to select Tools⇨Add a Contact — you can also select Tools⇨Search for Skype Users. Then (of course, your hardware has to be set up) you can click their name in the list and click the Phone button to give them a call!



The Skype network is a bandwidth-sharing application. After you connect to a Skype server, others can connect through your client by using your bandwidth. If you are on a limited-bandwidth connection, this could be a problem.



Figure 10-18:
The Skype
Internet
telephony
client.

Working with Other Internet Tools

You may want to explore a variety of other types of tools, depending on your needs. This section covers some pointers to get you started, so if you go to Chapter 16 and want to add the appropriate software (or find yourself digging

through menus and wondering what the program may be called), you aren't completely lost.

- ✓ **IRC programs:** Although the IM clients support IRC (Internet Relay Chat) these days, you may prefer a program that's used only for IRC. Common programs for this purpose are X-Chat and KIRC.
- ✓ **File sharing:** LimeWire is a popular file-sharing application if you want to use the Gnutella network.
- ✓ **RSS readers:** Firefox has a number of RSS extensions designed for it. In addition, Straw is a popular RSS program.

In general, if you're not sure what programs to use, do a Web search on the type of thing you want to do (like RSS) and the word *Linux*.

Chapter 11

Putting the X in Text

In This Chapter

- ▶ Viewing the contents of text files
 - ▶ Manipulating text files in `nano`
 - ▶ Working with text files in `gedit`
-

From text editors to word processors, Linux offers a wide variety of options for working with words. In this chapter, I take a look at different ways to view the contents of a text file, using some simple text editors in both the non-GUI and GUI environments. In Chapter 12, I take a look at office suites for those who would rather do word processing.

Viewing the Contents of a Text File

Almost all configuration files in Linux are text files. In addition, many pseudo-programs (called *shell scripts*), all HTML documentation, and many other items in your system are text files. Fortunately, if you just want to see what's in a text file and don't want to do anything to its contents, you don't have to use an editor or word processor. You can use three command-line commands to view text files: `cat`, `less`, and `more`. I bet that you will grow to love them.

Yes, that first command is `cat`, and it's taken from the word *concatenate*, which means to bring together end to end — you can use the `cat` command on multiple-text files to have their text joined, one file's contents directly after another's. Typically you use this command in the Linux world in the format `cat filename`, where the contents of the file `filename` are displayed on the screen. For example, if you create the short text file `greetings` and then type `cat greetings`, you see the following:

```
$ cat greetings
These are the contents of the greetings file.
Meow!
$ _
```

Of course, if the file contains more than a screen's worth of information, `cat` spews it all out at one time like a big hairball, and all but the last screen of text scrolls off the screen. It's a good thing that you have some other choices. The one you're likely to choose is `less`, which displays the contents of a file a full screen at a time. Then you press the spacebar to continue to the next screen. You can also use the arrow keys to move up and down one line at a time, if you want.

An alternative to `less` is `more`. The main difference between the two is that with `more`, you can move only forward through the file and see only a screen's worth of information at a time. You can't back up.

To use either `less` or `more`, the format is similar to the format used with the `cat` command: `less filename` or `more filename`. When you finish reading the document, press `Q` to exit.

Editing Text Files with nano

If you aren't using (or can't use) the GUI, you have numerous text editors available to you. The most powerful of these are `vi` and `emacs`, which you'll hear many people go on about. However, both these programs require a learning curve. For beginners who want to just edit the dang file and move on, I recommend `nano`, which offers a friendly, menu-driven interface.

To open a file in `nano`, type **`nano filename`**, such as **`nano file1`**. This action opens the file in the `nano` editor, as shown in Figure 11-1.

You can then edit or type in that file as much as you need to.

Saving your work as you go

To save your file's contents without closing it (so that you can keep working on it):

1. Press Ctrl+O to Write Out.

A prompt appears toward the bottom of the screen, asking for the name of the file and offering the current name as the default option.

Figure 11-1:
The file
file1
open in the
nano editor
in Fedora
Core.



2. If you want to use the same name, press Enter. If you want to change the name, make your changes and then press Enter.

The lower part of your screen now displays that it wrote (saved) a certain number of lines.

3. Get back to work!

Saving and moving on with your life

To save your file's contents and close it (because you're done working on it for now) — or to close the file and not save the changes — do the following:

1. Press Ctrl+X to Exit.

A prompt appears toward the bottom of the screen, asking whether it should save the *modified buffer*. This question is a fancy way of asking whether you want to save your changes.

2. Press Y to save your changes or press N to not save them.

If you press Y, you're asked for the name of the file and offered the current name as the default option. If you press N, nano closes and exits, so you're finished with these steps.

3. If you want to use the same name, press Enter. If you want to change the name, make your changes and then press Enter.

The nano editor closes, and your file is saved.

Going with *gedit*

You're not stuck with just command-prompt-based text editors in Linux. Lots of graphical options are available. In this section, I cover *gedit* (Figure 11-2) because it is the default GUI text editor for Fedora, which is on the DVD included with this book.




Figure 11-2:
The *gedit*
window
with a blank
file, in
Fedora.

Entering and editing text in *gedit*

gedit is strictly a *text editor*, in that you use it to generate raw text, whereas a *word processor* creates marked-up text that can be opened only by programs that can read that word processor's file formatting. If you want to add bold, italics, underlines, or any other special features to your document, proceed to Chapter 12.

To enter text in *gedit*, just click within the big white space and start typing. You have access to the standard collection of editing tools, such as cut, paste, and copy. To use these, select the text you want to work with and then click the appropriate button on the *gedit* toolbar (or right-click and choose the appropriate command from the context menu).

The really interesting thing about this particular text editor is its plug-ins. To use these features, you need to follow these steps:

1. Choose Edit⇨Preferences in *gedit*.

This action opens the Preferences dialog box.

2. Click the Plug-Ins tab.

The Plug-Ins tab's contents appear.

3. Click an item you're interested in within the Plug-Ins tab.**4. Click the About Plug-In button to get more information.**

The information is contained within the small About window that appears.

5. Click Close to get rid of the About window.**6. If you want to use this plug-in, click in its check box.**

The plug-in is activated if a check appears in the check box.

7. If the Configure Plug-In button becomes active for the plug-in you just selected, click the button to open the tool's plug-in configuration dialog box.

This dialog box will be different depending on which plug-in you're using.

8. When you're finished with the individual plug-in's configuration, click OK to return to the Preferences dialog box.**9. If you want to examine more plug-ins, return to Step 3.****10. When you're finished selecting plug-ins, click Close to close the Preferences dialog box.**

You can now access the plug-ins from your `gedit` menus. Each one is placed in its appropriate location: for example, Change Case appears on the Edit menu.

Saving your work

As with most programs, you have two choices for saving your work. You can save your work and keep going or save it and then close the program. To just save the file and keep going, follow these steps:

1. Click the Save button.

This button looks like a floppy disk. If you have never saved this file, clicking it opens the Save As dialog box.

2. Either choose one of the folders in the Save In Folder drop-down list box to save your file into, or click the right-facing arrow next to Browse for Other Folders to open the file browsing section.

The dialog box expands to allow you to browse to the location you want to use.

- 3. If necessary, browse through the directories in the left or right pane until you're in the directory where you want to save the file.**

Double-click the name of a directory to enter it or double-click the entry in the left pane to jump to its location first.

- 4. Type the file's name in the Name text box.**

- 5. Click Save to save the file.**

The dialog box closes.

To close `gedit`, follow these steps:

- 1. Choose File⇨Quit.**

If you haven't saved this file since the last time you changed it, the Question dialog box appears.

- 2. If you see the dialog box, click Save to save your work or click Close Without Saving to abandon it.**

The program closes, unless you have more than one file open, in which case you see the Question dialog box for each file you have altered but not saved.

Chapter 12

Word Processing and More with OpenOffice.org

In This Chapter

- ▶ Installing the OpenOffice.org suite
 - ▶ Word-processing with OpenOffice.org Writer
 - ▶ “Spreadsheets” with OpenOffice.org Calc
 - ▶ Presenting with OpenOffice.org Impress
 - ▶ Drawing with OpenOffice.org Draw
 - ▶ Managing data with OpenOffice.org Base
 - ▶ Formulating with OpenOffice.org Math
-

Words fly, writing remains.

— *Spanish proverb, from Dictionary of Proverbs, by Delfin Carbonell Basset*

These days, just about everyone who has a computer has at least one office suite at their fingertips. If they’re Microsoft Windows users, this suite is probably Microsoft Office, although it may be another worthy contender, such as Corel WordPerfect Office. In Linux, typically the suite is OpenOffice.org. This suite comes with Base (database), Calc (a spreadsheet), Draw (diagrams and figures), Impress (for presentations), Math (a word processor for writing mathematical formulas), and Writer (for word processing).

After you figure out how to use one of the programs in this suite, you may be happy to find that the others are designed to look and work in very similar ways. You can even open and save files in Microsoft Office format, if you need to share them with people who are using it — and you can edit the Office files people send you, too.

That’s enough *about* OpenOffice.org. In this chapter, you can actually *use* it!



Other office suites available for Linux users are the commercial OpenOffice.org relative StarOffice (www.sun.com/staroffice), Applixware Office (www.vistasource.com/products), KOffice (www.koffice.org), and GNOME Office (www.gnome.org/gnome-office/).

Installing the OpenOffice.org Suite

Before you can take off creating your masterpieces, you must make sure that OpenOffice.org is available. When you installed Fedora, one of the installation options was to install the Office productivity software. If you selected this option, you installed most of the OpenOffice.org suite of programs.

By default, Fedora includes these OpenOffice.org packages in the Office productivity installation:

- ✓ **Writer** for word processing
- ✓ **Calc** for spreadsheets
- ✓ **Impress** for presentation graphics
- ✓ **Draw** for drawing graphics
- ✓ **Math** for creating complex mathematical formulas

One additional OpenOffice.org package is available that is not installed by default: Base, which is used for creating personal databases.

If you chose not to install the Office productivity software, or you need to install the Base database software, don't sweat it, you can do that now without having to completely reinstall Fedora. Just use Chapter 16 to find how to add the OpenOffice.org packages from the Applications⇨Office/Productivity section's optional packages.

After you've installed the OpenOffice.org packages, you'll find most of them listed in the Applications⇨Office menu area (even the additional Base package if you install it). The two oddballs are the Draw and Math packages, which unfortunately don't rate to get a listing of their own.

To access either of these packages, you'll need to start one of the other OpenOffice.org packages (such as Writer), and then select File⇨New. In the New menu area you'll see the options for creating a new Drawing (which starts Draw) or a new Formula (which starts Math).



The Writer, Calc, and Impress packages even have their own quick-launch icons on the top panel. Just a single click of the icon launches the desired package.

Word Processing with OpenOffice.org Writer

Word processors are almost required equipment these days. Kids use them to write letters to their grandparents. Grandparents use them to write letters to their grandkids. Whether you're working on the great American novel or a school book report, OpenOffice.org Writer has all the best features you expect to find these days in a word processor.

Taking a tour of OpenOffice.org Writer

Before you proceed, take a look at the GUI layout shown in Figure 12-1.


Menu bar

Along the top of the window is the menu bar, something you should be used to if you typically work in Microsoft Windows. OpenOffice.org Writer has all the features you expect from a modern word processor. It has too many menu options to cover in depth, so I give you instead a (nonexhaustive) summary of what you find on each major menu:

- ✓ **File:** The usual Open, Save, Save As, Print, and Print Preview (under the term Page Preview) commands, along with a set of wizards (under the term AutoPilot), plus the ability to send documents through e-mail, create templates, create Web pages, and to access your database information
- ✓ **Edit:** The usual Select All and Find commands, along with change tracking, document merging, and document comparing
- ✓ **View:** The usual Zoom functions and toolbars, along with the abilities to show or hide formatting characters, and to see what the document would look like as a Web page
- ✓ **Insert:** The usual page breaks and special characters, along with indexes, tables, bookmarks, headers, footers, and cross-references
- ✓ **Format:** The usual character, paragraph, and page settings, along with styles, autoformatting capabilities, and columns
- ✓ **Table:** The usual table control options, including the ability to insert, delete, and select cells, convert between tables and text, and more
- ✓ **Tools:** The usual spell-checking and thesaurus entries, in addition to hyphenation, autocorrection, an image gallery, and a bibliography database

These menus have more features than what is listed here. Go through and take a look; you may find a new favorite feature in there somewhere.

Figure 12-1:
The
OpenOffice.
org Writer
layout in
Fedora.



Standard toolbar

Beneath the menu bar is the standard toolbar. Each icon in this series represents a different functionality, which is detailed in Table 12-1.

Table 12-1 **The OpenOffice.org Writer Standard Toolbar,
from Left to Right**

<i>Button or Item</i>	<i>What You Can Do</i>
New	Create new documents of various types. Click the downward arrow to select a particular type of document to create, from among any of the OpenOffice.org (OOo) types.
Open	Open an existing file for reading or editing.
Save	Save the current document. If you haven't saved this document before, the Save As dialog box opens.
Document as E-mail	Open a Compose e-mail window in your preferred e-mail program and automatically attach this document.
Edit File	Edit the displayed Web page.
Export Directly as PDF	Open a Save As dialog box with PDF selected as the file type.
Print File Directly	Send a file to the default printer.

<i>Button or Item</i>	<i>What You Can Do</i>
Page Preview	Show this page as it would look if you printed it. To return from preview mode, click Close Preview.
Spellcheck	Run the spell checker on your entire document or the selected text.
AutoSpellcheck On/Off	Activate or turn off the automatic spell checker feature.
Cut	Remove the selected text from the document and save it in memory.
Copy	Make a copy of the selected document text and save it in memory.
Paste	Place the text from memory into the document at the cursor's current location. Click the down arrow to see options of how the text can be pasted.
Format Paintbrush	Pick up the formatting of the first text you click and apply it to the second text you click.
Undo	Undo the last change you made to the document. Click the down arrow to choose how far you want to back up to.
Redo	Reinstate the last change to the document after using Undo to cancel it. Click the down arrow to choose how far you want to redo.
Hyperlink	Open or close a dialog box that you can use to build complex hyperlinks.
Table	Insert a new table. Click the down arrow to drag and choose how many rows and columns the table should have.
Show Draw Functions	Access the many OpenOffice.org drawing utilities.
Find and Replace	Open or close the Find and Replace dialog box.
Navigator	Open or close the Navigator window, which allows you to jump to specific features within your document.
Gallery	Open or close a pane along the top of the document that provides access to clip art. Click this button again to close the pane.
Nonprinting Characters On/Off	Shows all spaces, return at the ends of paragraphs, and other characters that you don't normally see in your documents.
Zoom	Alter how large the document shows on your screen.
Help	Open the OpenOffice.org Help dialog box.

Formatting toolbar

The Formatting toolbar is directly below the standard toolbar in a default OpenOffice.org setup — though not on all distributions. As usual, you can remove the Formatting bar at any time by using the View menu. This series of icons allows you to click buttons and expand drop-down list boxes that represent standard word-processing functions, such as styles, fonts, font sizes, and formatting instructions. Most features on this bar are identical to what you see in most modern word processors. The button for paragraph background formatting is the only one that's particularly unusual.



This toolbar actually changes depending on what you're doing. If the cursor is within a table, for example, the object bar contains useful buttons for working with tables.

Ruler

Directly below the Formatting toolbar in a default OpenOffice.org setup is the ruler. All modern word processors offer this item, which marks out the margins and tabs, for example, of your document in the measuring system of your choice. To change which system you want to use, right-click the ruler to open the Measurements pop-up dialog box.

Your document

Oh, yeah — that big, blocked-off white space that takes up most of the window. That's where you work on your documents! Just click in there and start typing. You can also access a Formatting shortcut menu by right-clicking in the document section.

Working with Writer files

OpenOffice.org Writer can work with lots of different file types. This allows you to create documents that can be used with many different types of word processing systems.

Not only can you save new documents in various formats, but you can also read documents created on other word processing programs. Here's a list of the different document types you can use with OpenOffice.org Writer:

- ✓ **OpenDocument (.ODT) format:** A proposed open standard that may one day be supported by all word processing programs.
- ✓ **OpenOffice.org 1.0 (.SXW) format:** The original OpenOffice document format.
- ✓ **Microsoft Word 97/2000/XP (.DOC) format:** Yes, you can read and write Word documents with Writer.

- ✓ **Microsoft Word 95 (.DOC) format:** And even some of the older Word formats.
- ✓ **Microsoft Word 6.0 (.DOC) format:** Even some of the really old Word formats.
- ✓ **Rich Text (.RTF) format:** A standard format for saving basic font and document information.
- ✓ **StarWriter 3.0 – 5.0 (.SDW) format:** Sun’s StarOffice word processor (a close cousin to OpenOffice.org).
- ✓ **Microsoft Word 2003 (.XML) format:** An attempt by Microsoft to create a standard word processing document format by using the XML standard language.

Although not available in the Save As dialog box, the OpenOffice.org Writer toolbar allows you to export any format document to the PDF document format. This feature alone can save you from having to purchase an expensive commercial product (from an unnamed commercial software company).

Spreadsheets with OpenOffice.org Calc

Some people like to balance their checkbooks by hand. When I first graduated from a university, I decided that it was time to get hold of my finances, and a spreadsheet was the way to do it. These days, I use spreadsheets to keep track of my “time card” when I’m doing consulting or contract work, help me manage project teams, and complete other tasks. I’m sure that you have your favorite uses for spreadsheets. The following sections take a look at OpenOffice.org Calc so that you can get to work.


Taking a tour of OpenOffice.org Calc

Much of what you see in OpenOffice.org Calc should look familiar, between looking through OpenOffice.org Writer and other spreadsheet programs you have used. Take a look at the GUI layout shown in Figure 12-2.

Menu bar

Along the top of the window is the menu bar, a standard in the GUI world no matter which operating system you’re using. OpenOffice.org Calc has all the features you expect from a modern spreadsheet system. It has too many menu options to cover in depth, so, instead, here’s a (nonexhaustive) summary of what you find on each menu:

Figure 12-2:
The
OpenOffice.
org Calc
layout in
Fedora.



- ✓ **File:** The usual Open, Save, Save As, Print, and Print Preview (under the term Page Preview) commands, along with a set of Wizards; plus the ability to send documents through e-mail, create templates, and create Web pages
- ✓ **Edit:** The usual Select All and Find commands, along with change tracking, headers and footers, and plug-in loading
- ✓ **View:** The usual Zoom functions and toolbars, along with the options for showing or hiding column and row headers
- ✓ **Insert:** The usual page breaks and special characters, along with cells, rows, functions, and external data
- ✓ **Format:** The usual cell and row formatting, cell merging, and page settings, along with conditional formatting
- ✓ **Tools:** The usual spell-checking and thesaurus entries, in addition to hyphenation, autocorrection features, an image gallery, and a macro creator and editor
- ✓ **Data:** The usual data selection, sorting, and grouping routines in one easy place for quick access

These menus have more features than those listed here. Go through and take a look; you may find a new favorite feature in there somewhere.

Standard toolbar

Beneath the menu bar is the standard toolbar. Each icon in this series represents a different functionality, as shown in Table 12-2.

Table 12-2 The OpenOffice.org Calc Standard Toolbar, from Left to Right	
<i>Button or Item</i>	<i>What You Can Do</i>
New	Open new documents of various types. Click the downward arrow to select a particular type of document to create, from among any of the OOo types.
Open	Open an existing file for reading or editing.
Save	Save the current document. If you haven't saved this document before, the Save As dialog box opens.
Document as E-mail	Open a Compose e-mail window in your preferred e-mail program, and automatically attach this document.
Edit File	Edit the displayed spreadsheet.
Export Directly as PDF	Open a Save As dialog box with PDF selected as the file type.
Print File Directly	Send a file to the default printer.
Page Preview	Show this page as it would look if you printed it. To return from preview mode, click Close Preview.
Spellcheck	Run the spell checker on your entire document or the selected text.
AutoSpellcheck	Activate or turn off the automatic spell checker feature.
Cut	Remove the selected text from the document and save it in memory.
Copy	Make a copy of the selected document text and save it in memory.
Paste	Place the text from memory into the document at the cursor's current location. Click the down arrow to see options of how the text can be pasted.
Format Paintbrush	Pick up the formatting of the first text you click and apply it to the second text you click.

(continued)

Table 12-2 (continued)

<i>Button or Item</i>	<i>What You Can Do</i>
Undo	Undo the last change you made to the document. Click the down arrow to choose how far you want to back up to.
Redo	Reinstate the last change to the document after using Undo to cancel it. Click the down arrow to choose how far you want to redo.
Hyperlink	Open or close a dialog box that you can use to build complex hyperlinks.
Sort Ascending	Re-order the selected data in ascending order.
Sort Descending	Re-order the selected data in descending order.
Insert Chart	Create a chart based on the selected data.
Show Draw Functions	Access the many OpenOffice.org drawing utilities.
Find and Replace	Open or close the Find and Replace dialog box.
Navigator	Open or close the Navigator window, which allows you to jump to specific features within your document.
Gallery	Open or close a dialog box that provides access to clip art.
Zoom	Alter how large the document shows on your screen.
Help	Open the OpenOffice.org Help dialog box.

Formatting bar

The object bar is directly below the function bar in a default OpenOffice.org setup. As usual, you can remove the object bar at any time by using the View menu. This series of icons allows you to click buttons and expand drop-down list boxes that represent standard spreadsheet functions, such as styles, fonts, font sizes, and number formatting instructions. Most features on this bar are identical to what you see in most modern spreadsheets.

Formula bar

Directly below the object bar in a default OpenOffice.org Calc setup is the Formula bar. Table 12-3 lays out what you find in this short collection of entries. This bar actually changes depending on what you're doing, offering you buttons for particular tasks, so don't panic if you look here and this table doesn't match what you see on your own Formula bar.

Table 12-3 **The OpenOffice.org Calc Formula Bar, from Left to Right**

<i>Button or Item</i>	<i>What You Can Do</i>
Function Wizard	Click to open the Function Wizard dialog box and browse to find the particular spreadsheet function you're looking for.
Sum	Click to start a SUM (addition) function in the Input Line.
Function	Click to place an = in the Input Line to signal that you're about to enter a function.
Input Line	Assign values or enter functions to fill a spreadsheet cell.

Your document

The document area is where you work on your spreadsheet. Just pick a cell and start typing. You can also access a Formatting shortcut menu by right-clicking in the document section.

Working with Calc files

Just like the Writer package, the Calc package allows you to read and write many different spreadsheet formats:

- ✓ **OpenDocument Spreadsheet (.ODS) format:** A proposed standard spreadsheet document format
- ✓ **OpenOffice.org 1.0 Spreadsheet (.OTS) format:** The older OpenOffice 1.0 spreadsheet format
- ✓ **Data Interchange Format (.DIF):** A text file format used to import and export spreadsheets between dissimilar spreadsheet programs
- ✓ **dBase (.DBF):** Exports the spreadsheet data into a dBase database file
- ✓ **Microsoft Excel 97/2000/XP (.XLS) format:** The standard Microsoft Excel spreadsheet formats used by the most popular versions of Excel
- ✓ **Microsoft Excel 95 (.XLS) format:** And even the older Excel spreadsheet format
- ✓ **Microsoft Excel 5.0 (.XLS) format:** And older Excel spreadsheet formats
- ✓ **StarCalc 3.0 – 5.0 (.SDC) format:** Sun's StarOffice spreadsheet format (a close cousin of the OpenOffice.org suite)
- ✓ **SYLK (.SLK) format:** The Symbolic Link format, used to exchange data between spreadsheets and other applications, such as databases

- ✓ **Text CSV (.CSV) format:** The comma separated text values, often used to export data to databases
- ✓ **HTML (.HTML) format:** Formats spreadsheet data as an HTML Web page
- ✓ **Microsoft Excel 2003 (.XML) format:** A standard proposed by Microsoft for using the Internet XML standard for defining spreadsheet data

Also, similar to Writer, the Calc standard toolbar provides a button for you to export the spreadsheet directly as a PDF document.

Presentations with OpenOffice.org Impress

Most people would rather eat glass than speak in front of a group. Still, if you have to, you may as well have some cool presentation software to back you up. This program is what I use when I speak at conferences. (I use it even when my clients want me to send in my presentation in Microsoft PowerPoint format. Shh! Don't tell — they don't know the difference!) Give OpenOffice.org Impress a chance to impress you.

Using the Presentation Wizard

When you open OpenOffice.org Impress, the first thing that launches is the Presentation Wizard (see Figure 12-3), which you can also open later manually by choosing **File**→**Wizards**→**Presentation**. If you don't want this wizard to show up the next time you open OpenOffice.org Impress, then click the **Do Not Show This Dialog Again** box.



Figure 12-3:
The
OpenOffice.
org
AutoPilot
Presenta-
tion Wizard
in Fedora.

To use this wizard, follow these steps:

1. **Leave the Empty Presentation option selected (unless you have a template you need to work from) and click Next.**

Dialog 2 appears.

2. **If you want to select one of the default slide backgrounds that come with OpenOffice.org Impress, you can do so by looking to the Select a Slide Design drop-down list box, choosing Presentation Backgrounds, and then clicking the various options to see what they look like to the right of the dialog box.**

By default, there aren't many background templates. This is one area where Sun's StarOffice offers added value.

3. **If you want to select a presentation (content) layout template that was designed specifically for the template you chose in this dialog box's screen 1, you can do so by looking to the Select a Slide Design drop-down list box, choosing Presentation(s), and then selecting the presentation type you want to use.**

Again, you find limited options here by default.

4. **If you want your presentation ultimately to appear on something other than a computer screen, adjust the Select An Output medium to match its intended setting.**

Your choices are Screen, Overhead Sheet, Slide, and Paper.

5. **Click Next to proceed.**

Dialog 3 appears.

6. **Under Select a Slide Transition, experiment with the various options in the Effect and Speed drop-down list boxes to narrow down how you want to move from one slide to another.**

OpenOffice.org Impress animates these transitions for you as long as the Preview box is checked.

7. **If you want to navigate manually from one slide to the next (the Default option) while you give your presentation, skip to Step 10. If you want to have your presentation advance automatically, click the Automatic option.**

8. **In Duration of Page, set how long you want each slide to stay up.**

9. **In Duration of Pause, set how long of a blank gap you want to have between slides.**

If you have Show Logo checked, the OpenOffice.org logo appears during the blank pauses.

10. **After you have your settings selected, click Create to proceed.**

OpenOffice.org Impress opens, as shown in Figure 12-4. You may find it useful to click the X in the upper right of the Slides pane to clear up the window some.




Figure 12-4:
The
OpenOffice.
org Impress
default look
in Fedora.

Taking a tour of OpenOffice.org Impress

Before you proceed, take a look at the GUI layout shown in Figure 12-4. The Slides view on the left and the Tasks view on the right might look like clutter to some; if you don't need them, just click the X in the upper-right corner of each pane. You can bring them back at any time by choosing View⇨Slide Pane and/or View⇨Task Pane.

Menu bar

Along the top of the window is the menu bar. OpenOffice.org Impress has the many features you expect from a modern presentation package. It has too many menu options to cover in depth, so I give you instead a (nonexhaustive) summary of what you find on each menu:

- ✓ **File:** The usual Open, Save, Save As, and Print commands, along with a set of Wizards, plus the ability to send documents through e-mail and create templates
- ✓ **Edit:** The usual Select All and Find commands, along with the ability to quickly duplicate a slide

- ✓ **View:** The usual Zoom functions and toolbars, along with the ability to select whether you're looking at just slides, notes, or another section
- ✓ **Insert:** The usual new slide, along with charts, frames, graphics, and spreadsheets
- ✓ **Format:** The usual text-formatting features, along with layout, graphics, and style formatting
- ✓ **Tools:** The usual spell-checking feature, in addition to hyphenation, auto-correction, and an image gallery
- ✓ **Slide Show:** The usual slide show controller menu



These menus have more features than those listed here. Don't forget to do some exploring on your own.

Standard toolbar

Along the top of the window is the standard toolbar, which you can remove at any time by choosing View⇨Toolbars⇨Standard. Each icon in this series represents a different functionality and is described in Table 12-4.

Table 12-4 **The OpenOffice.org Impress Standard Toolbar, from Left to Right**

<i>Button or Item</i>	<i>What You Can Do</i>
New	Open new documents of various types. Click the downward arrow to select a particular type of document to create, from among any of the OOo types.
Open	Open an existing file for reading or editing.
Save	Save the current document. If you haven't saved this document before, the Save As dialog box opens.
Document as E-mail	Open a Compose e-mail window in your preferred e-mail program and automatically attach this document.
Edit File	Edit the displayed file.
Export Directly as PDF	Open a Save As dialog box with PDF selected as the file type.
Print File Directly	Send a file to the default printer.
Spellcheck	Run the spell checker on your entire document or the selected text.
AutoSpellcheck	Activate or turn off the automatic spell checker feature.

(continued)

Table 12-4 (continued)

<i>Button or Item</i>	<i>What You Can Do</i>
Cut	Remove the selected text from the document and save it in memory.
Copy	Make a copy of the selected document text and save it in memory.
Paste	Place the text from memory into the document at the cursor's current location. Click the down arrow to see options of how the text can be pasted.
Format Paintbrush	Pick up the formatting of the first text you click and apply it to the next text you click.
Undo	Undo the last change you made to the document. Click the down arrow to choose how far you want to back up to.
Restore	Reinstate the last change to the document after using Undo to cancel it. Click the down arrow to choose how far you want to redo.
Chart	Insert a chart into the presentation by using the selected data.
Spreadsheet	Insert a spreadsheet into the presentation.
Hyperlink	Open or close a dialog box that you can use to build complex hyperlinks.
Display Grid	Display or remove the line-up grid from the slide.
Navigator	Open or close the Navigator window, which allows you to jump to specific features within your document.
Zoom	Alter how large the document shows on your screen.
Help	Open the OpenOffice.org Help dialog box.
Slide	Insert a new slide after the current one.
Slide Design	Open the Slide Design dialog box.
Slide Show	Start a slide show.

Line and filling bar

The line and filling toolbar is directly below the standard toolbar in a default OpenOffice.org setup. As usual, you can remove the object bar at any time by using the View menu. This series of icons allows you to click buttons and expand drop-down list boxes that represent standard presentation-software functions, such as arrow styles, colors, line styles, and other formatting

instructions. Most features on this bar are identical to what you see in most modern presentation programs.

View switchers

Just above your document you see a series of tabs, each of which takes you to a particular way of viewing your slide(s). Table 12-5 outlines the available views and what you find in them.

Table 12-5 Available OpenOffice.org Impress Views, in Alphabetical Order

<i>View</i>	<i>What You Find</i>
Normal	Individual slide view in which you can add art to your slide.
Handout	Six slides per page, as you might print it for handouts.
Master	Individual slide view in which you can apply master formatting that will apply to all slides. Access this view by selecting View⇨Master. You have the option in the Master submenu of choosing either Slide Master to view the master slide, Notes Master to view the master note page, or Master Elements to assign the information that is made available in the Master views.
Notes	Individual slide view in which you can see a small version of the slide plus your notes about that slide.
Outline	All-slides view with the slides listed in order for easy stepping through. Along the side, the slides are shown in thumbnail mode in a separate window as you navigate.
Slide Sorter	All-slides view with as many slides packed in as possible in columns and rows. Re-ordering slides is as simple as dragging them where you want to go and then dropping them.

Your document

The tabbed window contains your document. In Normal view, that document is your slide. You need to click one of the tools in the Drawing Bar (discussed in the next section) in order to be able to enter any content. To add more slides before or after this one, right-click over the document area and select Slide⇨New Slide. For each slide, you can use the layouts in the Tasks pane on the right to change their setup.

Drawing bar

Beneath the document window is a drawing toolbar, which allows you to select lines, arrows, shapes, and more for your presentation creation needs.

Most of these buttons have downward-pointing arrows, which allow you to see the full range of features they offer.

Working with Impress files

The OpenOffice.org Impress application doesn't support as many file types as its Writer and Calc siblings. Mostly this is due to a lack of standards in the presentation graphics environment.

The formats that Impress does support are

- ✓ **OpenDocument Presentation (.ODP) format:** A standard presentation graphics format proposed by OpenOffice.org.
- ✓ **OpenOffice.org 1.0 Presentation (.OTP) format:** The original OpenOffice.org presentation graphics file format.
- ✓ **Microsoft PowerPoint 97/2000/XP (.PPT) format:** Possibly the most popular presentation graphics tool used in business. You can use Impress to read and write most PowerPoint slide presentations.
- ✓ **OpenOffice.org 1.0 Drawing (.SXD) format:** An older rudimentary graphics drawing format.
- ✓ **StarDraw 3.0 – 5.0 (.SDA) format:** Draw's commercial cousin from Sun.
- ✓ **StarImpress 4.0 and 5.0 (.SDD) format:** Impress's commercial cousin from Sun.
- ✓ **OpenOffice.org Drawing (.ODG) format:** Rudimentary graphics drawing format.

Although Impress can work with OpenOffice.org Draw files, the default file format for Impress files is the .ODP format. As with the other packages in the OpenOffice.org suite, Impress also provides a toolbar icon for you to save your slides as .PDF format files.

Fine Art with OpenOffice.Org Draw

Whether you're an aspiring graphic artist or just need a tool that lets you generate simple graphics for use on their own, in a presentation, or elsewhere, OpenOffice.org Draw provides a host of drawing functions. If nothing else, it's a whole lot of fun to play with! Not everything in life has to be practical.

Taking a tour of OpenOffice.org Draw

Before you proceed, take a look at the GUI layout shown in Figure 12-5. If you find the Pages pane on the left too much clutter, click the X in its upper right to get rid of it. You can bring it back at any time by selecting View→Page Pane.




Figure 12-5:
The
OpenOffice.
org Draw
layout in
Fedora.

Menu bar

Along the top of the window is the usual menu bar. OpenOffice.org Draw is a typical “vector” graphics program, meaning that it relies on lines rather than dots or other techniques. See Chapter 14 for discussion of the software used for editing photographs and other heavy-detail work.

OpenOffice.org Draw has too many menu options to cover in depth, so I give you instead a (nonexhaustive) summary of what you find on each menu:

- ✓ **File:** The usual Open, Save, Save As, Print, and Export commands, along with a set of Wizards (under the term AutoPilot) plus the ability to send documents through e-mail and create templates
- ✓ **Edit:** The usual Find, Replace, and other commands, and the not so usual Image Map
- ✓ **View:** The usual Zoom functions and toolbars, along with the ability to select the display quality
- ✓ **Insert:** The usual charts, frames, graphics, and spreadsheets, along with scanning functions



- ✓ **Format:** The usual line and graphics formatting, along with layers and style formatting
- ✓ **Tools:** The usual spell checker, as well as hyphenation, autocorrection, an image gallery, and an eyedropper for grabbing colors
- ✓ **Modify:** Various options for altering the appearance of an object

These menus have more features than those listed here. Go through and take a look; you may find a new favorite feature in there somewhere.

Standard toolbar

Beneath the menu bar is the standard toolbar, which you can remove at any time by choosing View⇨Toolbars⇨Standard. Each icon in this series represents a different functionality, as described in Table 12-6. As you can see, this main toolbar is more similar to that in OpenOffice.org Impress than to the one in OpenOffice.org Writer.

Table 12-6 **The OpenOffice.org Draw Main Toolbar,
from Left to Right**

<i>Button or Item</i>	<i>What You Can Do</i>
New	Open new documents of various types. Click the downward arrow to select a particular type of document to create, from among any of the OOo types.
Open	Open an existing file for reading or editing.
Save	Save the current document. If you haven't saved this document before, the Save As dialog box opens.
Document as E-mail	Open a Compose e-mail window in your preferred e-mail program, and automatically attach this document.
Edit File	Edit the document.
Export Directly as PDF	Open a Save As dialog box with PDF selected as the file type.
Print File Directly	Send a file to the default printer.
Spellcheck	Run the spell checker on your entire document or the selected text.
AutoSpellcheck	Activate or turn off the automatic spell checker feature.
Cut	Remove the selected text from the document and save it in memory.
Copy	Make a copy of the selected document text and save it in memory.

<i>Button or Item</i>	<i>What You Can Do</i>
Paste	Place the text from memory into the document at the cursor's current location. Click the down arrow to see options of how the text can be pasted.
Format Paintbrush	Pick up the formatting of the first item you click and apply it to the second item you click.
Undo	Undo the last change you made to the document. Click the down arrow to choose how far you want to back up to.
Restore	Reinstate the last change to the document after using Undo to cancel it. Click the down arrow to choose how far you want to redo.
Chart	Insert a chart using the selected data.
Hyperlink	Open or close a dialog box that you can use to build complex hyperlinks.
Navigator	Open or close the Navigator window, which allows you to jump to specific features within your document.
Zoom	Alter how large the document shows on your screen.
Help	Open the OpenOffice.org Help dialog box.

Line and filling bar

The line and filling bar is directly below the standard bar in a default OpenOffice.org setup. As usual, you can remove the line and filling bar at any time by using the View menu. This series of icons allows you to click buttons and expand drop-down list boxes that represent standard presentation software functions, such as arrow styles, colors, line styles, and other formatting instructions.

Ruler

Directly below the line and filling bar in a default OpenOffice.org setup are the rulers. These items mark out the margins and tabs, for example, of your document in the measuring system of your choice. To change which system you want to use, right-click the ruler and change the Measurements to your preferences.

Your document

Click in that big white space and start doodling. You can also access a formatting pop-up menu by right-clicking in the document section.

If you prefer to deface other pictures, you can use the File⇨Open menu selection to select any standard graphics format file (such as .BMP or .JPG) to

work with. After the file is opened in Draw, you can use all the standard Draw tools available on the picture.

Drawing bar

Beneath the document window is a drawing toolbar, which allows you to select lines, arrows, shapes, and more for your drawing creation needs. Most of these buttons have downward-pointing arrows, which allow you to see the full range of features they offer.

Working with Draw files

At first glance, OpenOffice.org Draw appears to have the most limited file type support for saving when you use the Save As dialog box. From there, the only file formats you can save a document as are

- ✓ **OpenDocument Drawing (.ODG) format:** The OpenOffice.org proposed standard for graphics files
- ✓ **OpenOffice.org 1.0 Drawing (.SXD) format:** The older OpenOffice.org graphics format
- ✓ **StarDraw 3.0 (.SDD) format:** The older Sun StarDraw format
- ✓ **StarDraw 5.0 (.SDA) format:** The newer Sun StarDraw format

That's not much to work with, especially when dealing with picture files. Fortunately, Draw has a few tricks up its sleeve. You can save images to a format other than the standard formats by choosing File→Export.

This produces a host of supported graphics formats: BMP, EMF, EPS, GIF, JPEG, MET, PBM, PCT, PGM, PNG, PPM, RAS, SVG, SWF (Flash), TIFF, WMF, and XPM. That should satisfy any computer artist.

Managing Data with OpenOffice.org Base


Those who are used to using Microsoft Access, or who have to interact with a variety of database products, have in the past been left scratching their heads and wondering how best to handle these needs under Linux. Well, wonder no more! OpenOffice.org Base provides an interface to both Microsoft Access files and to other databases.

Getting Help from the Wizard

There are many different ways you can use this program. The essentials involve creating a database from scratch in a file on your system, working with Microsoft Access files, and working with databases that were created to use with many of the major database servers — even over the network. Because approaches abound and my space is limited, I focus on two of the more common choices: creating new database files, and opening existing files, which include Microsoft Access files.

The first interface you encounter when opening anything but a document you already worked with from the File→Recent Documents menu option is the Database Wizard (Figure 12-6).

Figure 12-6: The OpenOffice.org Base Database Wizard, used to open existing databases and create new ones.



Creating a new database file in the Database Wizard

In the Database Wizard dialog box, do the following to create a brand-new database file:

1. **Select the Create a New Database radio button.**
2. **Click Next.**

The Decide How To Proceed After Saving The Database dialog box appears, as shown in Figure 12-7.

3. **If you want to register this database as a data source in OpenOffice.org, leave Yes, Register the Database for Me selected; otherwise, select the No, Do Not Register the Database radio button.**

A registered database is accessible by all of your OpenOffice.org applications, rather than just OpenOffice.org Base.

Figure 12-7:
The
OpenOffice.
org Base
Database
Wizard's
Decide How
to Proceed
After
Saving the
Database
dialog box.



4. If you immediately want to open this file for editing, leave the **Open the Database for Editing** check box checked; otherwise, click this option to uncheck the box.
5. If you want to immediately create a new table in the database by using the **Table Wizard** (using the wizard is recommended), select the **Create Tables Using the Table Wizard** check box.

For now, it's best to leave it off so you can explore one step at a time.

6. Click *Finish*.

The **Save** dialog box opens.

7. Enter the name for your document in the *Name* text box.

8. Either select the directory you want to save the document to in the *Save In Folder* drop-down list box, or click the right-facing arrow next to *Browse for Other Folders* to navigate to where you want to save the document.

9. Click the *Save* button.


Whatever you specified should happen, happens. If you chose to immediately open the database, you see something equivalent to Figure 12-8.

Opening an existing file in the Database Wizard

If you have an already existing Microsoft Access file that you want to open in OpenOffice.org Base, do the following from the initial Database Wizard dialog box:

1. **Select *Open an Existing File*.**
2. **If the file is in the *Recently Used* drop-down list box, select the file there. Otherwise, click the *Open* button.**

Figure 12-8:
A database
open in
OpenOffice.
org Base.



If you clicked the Open button, the Open dialog box appears. If you selected an item from Recently Used, click Finish, and you're done opening the file.

3. **Navigate to the file you want to open, and select the file.**
4. **Click Open.**

The file is opened into a window similar to that shown in Figure 12-8.

Taking a tour of OpenOffice.org base

Before you proceed, refer to the layout shown in Figure 12-8. This is your main OpenOffice.org Base window after you have a database open to work with. Here's a look at what you'll find there.

Menu bar

Along the top of the window is the menu bar. OpenOffice.org Base gives you access to many of the features you may expect in a database interface:

- ✓ **File:** The usual Open, Save, and Save As commands, along with a set of wizards and the ability to send documents through e-mail
- ✓ **Edit:** The usual Copy and Paste commands, along with access to database properties, advanced settings, and the Form and Report Wizards
- ✓ **View:** Access to toolbars and database objects, along with preview and sort features

- ✓ **Insert:** The ability to create forms, queries, and more
- ✓ **Tools:** The ability to assign relationships between tables, filter tables, run SQL queries, create macros, and more

Standard toolbar

Beneath the menu bar is the standard toolbar, which you can remove at any time by choosing View⇨Toolbars⇨Standard. Each icon in this series represents a different functionality, as described in Table 12-7. As you can see, this main toolbar is more similar to that in OpenOffice.org Impress than to the one in OpenOffice.org Writer.

Table 12-7 The OpenOffice.org Base Standard Toolbar, from Left to Right

<i>Button or Item</i>	<i>What You Can Do</i>
New	Open new documents of various types. Click the downward arrow to select a particular type of document to create, from among any of the OOo types.
Open	Open an existing file for reading or editing.
Save	Save the current document. If you haven't saved this document before, the Save As dialog box opens.
Copy	Make a copy of the selected document text and save it in memory.
Paste	Place the text from memory into the document at the cursor's current location. Click the down arrow to see options of how the text can be pasted.
Sort Ascending	Sort the entries in the lower-right pane in alphabetical order.
Sort Descending	Sort the entries in the lower-right pane in reverse alphabetical order.
Form	Create a form.
Help	Open the OpenOffice.org Help dialog box.
Open Database Object	Open the selected item in the lower-right pane.
Edit	Open the design view for the selected item in the lower-right pane.
Delete	Delete the selected item in the lower-right pane.
Rename	Rename the selected item in the lower-right pane.

The Database pane

To the left of the main window is the Database pane. In this pane you see four icons: Tables, Queries, Forms, and Reports. Select these icons to determine what appears in the two rightmost panes — for example, to work with your tables, select the Tables icon.



A *query* allows you to ask complex questions regarding your database.

The Tasks pane

The upper-right pane is the Tasks pane. In here, you see what types of things you can do for a particular selection in the Database pane. Use the Description to the right to decrypt any of the terms that make your eyes cross.

The Tables/Queries/Forms/Reports pane

The lower-right pane shows the tables, queries, forms, or reports that already exist. You can open any of these to work in by double-clicking them in this pane.

Sitting down at the table

Tables are the core of the database system. They are what hold your data, and allow you to quickly access the information. It's a good idea to understand how tables are organized in Base, and how you can use them to hold and display your data.

To create a new table, you can either use the Base Table Wizard if you think you have a standard table, or you can just create your own table in Design View. Figure 12-9 shows an empty Table Design View dialog box.

Each row in the Table Design represents a data field in the table. Each field must have a unique Field Name and be assigned a Field Type. The Field Type defines the type of data present in the field (characters, numbers, dates, and so on). The Description area allows you to write a brief description describing the meaning of the data field.

Creating a new table is easy:

- 1. Enter a text name for the Field Name describing the data.**
- 2. Select a Field Type from the drop-down menu.**
- 3. Enter a brief description in the Description text area.**
- 4. If the field is a primary key, right-click on the left of the row and select Primary Key.**




Figure 12-9:
The Table
Design View
dialog box
in Base.

A table's primary key is a data field that can uniquely identify each record in the table. In the example in Figure 12-9, each employee is identified by a unique employee ID because more than one employee can have the same last name.

5. Enter the remaining rows to complete the data elements in the table.
6. Select **File**→**Save As** to save the table, using a unique Table name.

The new table now appears on the main Base window under the Tables area. Double-clicking the new table brings up the data viewing and entry window, as shown in Figure 12-10.

If you are familiar with Microsoft Access, you should feel right at home here. This window allows you to view the existing data records in the table, and easily add new data records.

Layout with OpenOffice.Org Math

There's nothing like trying to type a math or science report and having to either use multiple lines to show your equations (which never looks right!) or write them by hand. OpenOffice.org Math is a great solution to this problem. You can lay out your equations and more in this program and then insert them into any of your OpenOffice.org documents. Whether you're a middle-school student or a professional engineer, OpenOffice.org Math just may thrill you.

Figure 12-10:
The Data
Viewing and
Entry
window in
the Base
application.

empid	lastname	firstname	salary
1	Smith	Fred	10000
2	Jones	Janey	25000
3	White	Barbara	34200
4	Brown	John	56400

Many OpenOffice.org Math functions are different from what you're used to if you have looked at all the other OpenOffice.org programs. However, in many ways this program is less complex than some, thanks to its special-purpose nature. Keep in mind that it's not a calculation program. It's for laying out complex formulas on paper or the screen.

Before you proceed, take a look at the GUI layout shown in Figure 12-11.

Close the Selection dialog box by clicking the small X in the upper-right corner of its window for now, to keep things as uncluttered as possible. You can get it back at any time by choosing View⇨Selection.

Menu bar

Along the top of the window is the menu bar, a standard in the GUI world no matter which operating system you're using. OpenOffice.org Math may be unlike any program you may have used already, so I don't say anything about what you may expect to find there. I just give you a (nonexhaustive) summary of what you find on each menu:

- ✓ **File:** The usual Open, Save, Save As, and Print commands that you find in most GUI programs, along with a set of Wizards and the ability to send documents through e-mail




Figure 12-11:
The
OpenOffice.
org Math
layout.

- ✔ **Edit:** The usual Select All, Copy, and Paste commands, along with specialized commands for moving within the formula
- ✔ **View:** The usual Zoom functions and toolbars, along with screen update features and more
- ✔ **Format:** The usual font type, font size, spacing, and alignment features and more
- ✔ **Tools:** The usual Customize and Options entries for customizing the program's setup and behaviors, in addition to formula importing and access to the symbol catalog

Standard toolbar

Along the top of the window is the standard toolbar, which you can remove at any time by choosing View⇧Toolbars⇧Standard. Each icon in this series represents a different functionality. You're likely to find this main toolbar quite different from those in the other OpenOffice.org programs. Mostly, it's just smaller. Each icon is described in Table 12-8.

**Table 12-8 The OpenOffice.org Math Standard Toolbar,
from Left to Right**

<i>Button or Item</i>	<i>What You Can Do</i>
Zoom In	Show the image larger.
Zoom Out	Shrink the image.
1	Show the image at its actual size.
Show All	Show the whole formula in the largest size that will fit on the screen.
New	Open new documents of various types. Click the downward arrow to select a particular type of document to create, from among any of the OOo types.
Open	Open an existing file for reading or editing.
Save	Save the current document. If you haven't saved this document before, the Save As dialog box opens.
Document as E-mail	Open a Compose e-mail window in your preferred e-mail program, and automatically attach this document.
Export Directly as PDF	Open a Save As dialog box with PDF selected as the file type.
Print File Directly	Send a file to the default printer.
Cut	Remove the selected text from the document and save it in memory.
Copy	Make a copy of the selected document text and save it in memory.
Paste	Place the text from memory into the document at the cursor's current location.
Undo	Undo the last change you made to the document. Click the down arrow to choose how far you want to back up to.
Restore	Reinstate the last change to the document after using Undo to cancel it. Click the down arrow to choose how far you want to redo.
Help	Open the OpenOffice.org Help dialog box.
Update	Update the formula shown in the document window.
Formula Cursor	Turn on or shut off the formula cursor.
Catalog	Open the Symbols dialog box.

Document section

Things get tricky here if you've never used formula-editing software. You can't type anything in the main (upper) document window in OpenOffice.org Math. Instead, you type in the Commands (lower) window. Right-clicking in the Commands window opens a shortcut menu. To help you get used to working with the formulas, I suggest that you play around with this tool. For example, if you have never used software like this, you might follow these steps:

- 1. Right-click in the Commands window.**

The main shortcut menu opens.

- 2. Select a submenu to open.**

For example, *Formats*.

- 3. Select a formula component within this submenu.**

I'm a geek, and I fondly remember taking a class on matrices, so I chose matrix {...} as an example. Immediately, the code that's needed in order to add a matrix to my formula appears in the Commands dialog box. A moment later, because I'm letting the program refresh the rest of the screen as I work, I see what the matrix looks like in the document window. The combination is shown in Figure 12-12.




Figure 12-12:
Starting
to add a
matrix in
OpenOffice.
org Math.

4. Replace each of the `<?>` entries with the proper letters and numbers for your formula.

When I change `matrix{<?> # <?> ## <?> # <?>}` to `matrix{A # B ## C # D}`, I see the result shown in Figure 12-13.




Figure 12-13:
A 4-x-4
matrix in
OpenOffice.
org Math.

5. Continue adding components to the formula until you finish.

Suppose that you want to multiply the matrix by 3. To find out how, press Enter to go down to the next line in the Commands window (to use as “scratch paper”), right-click to display the pop-up menu, and choose Unary/Binary Operators → a Times b. This choice adds the phrase `<?> times <?>` beneath the matrix code. Now you know how to format a multiplication, so erase this phrase and use it as a guideline to change your formula to:

```
3 times matrix{A # B ## C # D}
```

This line gives the result shown in Figure 12-14. I could go on, but I hope by now that you’re eager to start with your own explorations!

OpenOffice.org Math supports (among others) its own OpenDocument format (.ODF), the equivalent StarOffice format (.SXM) and MathML 1.01 (.MML), which isn’t a program. *MathML* is a standard similar to HTML for

working on the Web, and 1.01 is a specific version of this standard. For this particular version of MathML, see www.w3.org/TR/REC-MathML. The main standard page is available at www.w3.org/Math/. If you need to add formulas to Web pages, this site can be quite an interesting read!

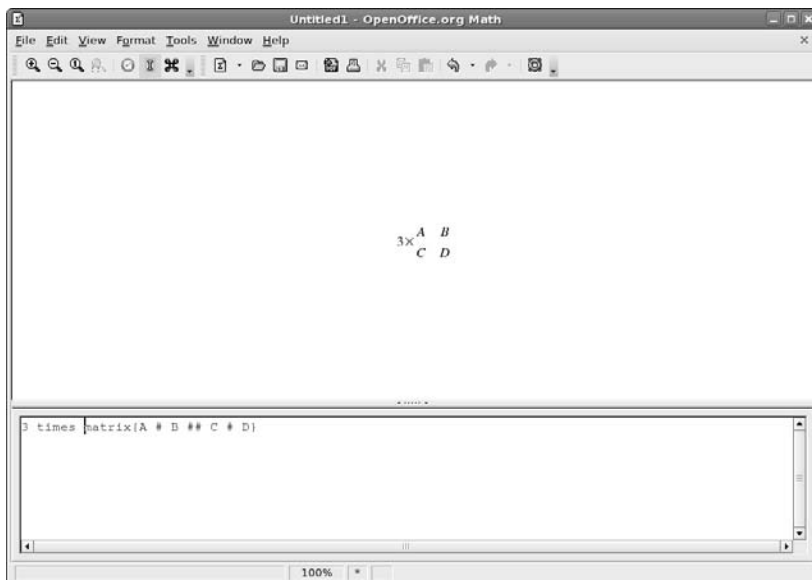


Figure 12-14:
A complete
formula in
OpenOffice.
org Math.

Chapter 13

Messing with Audio

In This Chapter

- ▶ Listening to music
 - ▶ Using, recording to, and ripping from CDs
 - ▶ Fixing sound problems
-

The price of freedom is responsibility, but it's a bargain, because freedom is priceless.

— Hugh Downs

If you're like me, everything's better with music. If a computer can't take the place of a stereo wherever I am, then I'm just not interested. So, I can assure you that there are many ways to listen to music within Linux. Read on to find the various methods available! The only tricky bit used to be listening to MP3s, because they're a patented format, but RealPlayer is available for Linux and has a license to play MP3s, so you're all set. See Chapter 16 for how to install RealPlayer.



Are you an iTunes fan? If so, check out Chapter 15 for how you can install iTunes in Linux!

What Sound? I Don't Hear a Thing!

First, check and make sure that your speakers or headphones are hooked up correctly. You wouldn't be the first person to be unaware that the cat pulled the connectors out of their sockets.

Next, it's possible that your audio settings got turned down to zero, or muted. Right-click the volume icon in the upper left of your desktop, and select Open Volume Control (Figure 13-1).




Figure 13-1:
The Fedora
Volume
Control.

When this control is open, it helps to understand the terms you're about to see:

- ✓ **Master:** The main volume control corresponds to the volume slider on your desktop.
- ✓ **PCM:** Controls the volume for audio files stored on your computer.
- ✓ **Line-in:** Controls the volume for sound coming in through your sound card's line-in socket.
- ✓ **CD:** Controls the volume for sound coming in through your CD and/or DVD player.
- ✓ **Microphone:** Controls the volume for sound coming in through your sound card's microphone connector.

Now that you know what you're looking at in this dialog box, look for one of these two major clues as to what might be wrong:

- ✓ Are any of the items you see muted? You can tell by looking for red X's on top of the volume icon underneath each section. In Figure 13-1, for example, Line-in and Microphone are both muted. Is the one you see muted responsible for your problem? For example, if you had headphones hooked up to your Line-in jack and Line-in was muted, you would put on your headphones and hear nothing. To unmute, click the X to get rid of it.

- ✓ Are any of the volume sliders all the way down? Again, Line-in and Microphone in Figure 13-1 are both slid down all the way. Using the same headphone example, even if the Line-in wasn't muted, you would hear nothing because the sliders are down. Slide them up and see if you can hear anything.

If the problem isn't one of these, then it may be something more fundamental. To investigate:

1. Choose System⇨Administration⇨Soundcard Detection (Figure 13-2) to see whether the card is correctly detected and plays test sounds.

You may be asked to enter your root (administrative) password before the tool will open.




Figure 13-2:
The Fedora
Audio
Configu-
ration tool.

2. If you know what kind of sound card you have, look under Selected Card.

Is the information correct? If not, click the Settings tab and click the drop-down list box next to Default Audio Card. If the proper choice is available, select it.

3. Click the play button under Sound Test.

You should hear a sound on the left, the right, and in the middle. If not, click the Repeat check box, click Play again, and then adjust the Volume Settings bar beneath to see if you can make the sound loud enough to hear.

4. If you still can't hear anything, change the PCM Device setting and try playing the sound again.

If this test doesn't work, change the PCM Device back to what it was originally before moving on.

5. If none of these actions helps, go to your favorite search engine and search on the make and model of your sound card, and the word **linux**.

You may find some helpful information that way.

Listening to CDs

By default, when you put a music CD into a computer running Fedora, the Totem movie/music player opens. However, this program doesn't automatically grab information about the CD, which personally I find pretty annoying. So, the first thing I recommend doing is switching CD players. To do so:

1. Select **System**⇨**Preferences**⇨**Hardware**⇨**Removable Drives and Media**.

The Removable Drives and Media Preferences dialog box appears.

2. Click the **Multimedia** tab.

The Multimedia section (Figure 13-3) appears.

Figure 13-3:
The
Removable
Drives and
Media
Preferences
dialog box
with the
Multimedia
tab
selected.



3. Under Audio CD Discs, replace totem with gnome-cd.

Leave the %d where it is.

4. Click Close.

The dialog box closes.

Now when you insert an audio CD, the CD player (Figure 13-4) opens instead of the movie player.

Figure 13-4:
The GNOME
CD Player,
with a CD
playing.




If the program doesn't recognize your CD, you can add it to the database by:

1. In the row above the bottom, click the Open Track Editor button.

The CDDB Track Editor opens. (See Figure 13-5).

Figure 13-5:
The CD
Player's
CDDB Track
Editor dialog
box in
Fedora.



2. Fill in the Artist, Disc Title, and the individual tracks.

To fill in a track, click the row the track belongs in, and then click the section you want to edit (such as the title). Those actions open an editable text section you can type in.

3. If you want to enter more information about the CD, click the arrow beside Show Advanced Disc Options, and fill in the additional text boxes.

This section allows you to enter the CD's genre and other such data.

4. If you want to enter extra information about one or more tracks, click the arrow next to Show Advanced Track Options, select the track you want to make notes on, and then make your notes in the Extra Track Data text box.

This text box is tied to the track you have selected, so when you select another track you'll find it empty.

5. When you finish, click Save.

On the left of the same row of controls, you'll find the CD Player Preferences button. Click it to open the window of the same name. There, you can set the preferences for your various CD/DVD players if you have more than one, select a theme for your CD player, and more.

Listening to Downloaded Music

Downloading music from the Internet is a fun activity. A lot of people like to pretend that it's both legal and ethical, when in fact it's often theft, depending on where you grabbed the music. I leave that ethical issue between you and your belief system, but because my own ethics say that it's theft, I'm going to focus on showing you music that you're *welcome* — legally and ethically — to download!

One place to find such material is on a band's own Web site. Many groups today realize that offering free downloads of their work is a great PR move — if people love the sample songs, they're more likely to go out and buy the album. For an example, I use the popular band Evanescence (www.evanescence.com).

Evanescence's Web site, like those of many other bands, is done mostly with Macromedia Flash. (See Chapter 16 for how to add this functionality to your distribution.) To listen to music files in general, I personally recommend Amarok (<http://amarok.kde.org/>), which (again) you can get by using

the tools discussed in Chapter 16 — search for `amarok` and install `amarok`, `amarokFS`, and `amarok-visualisation`. This feature-rich and professional-looking tool lets Linux desktop users feel like they're actually in the modern age. I don't have room to cover every option in exhaustive detail, but I'll give you the nickel tour at least.

1. **After you have Amarok installed, start it up by going to Applications⇨ Sound & Video⇨amaroK.**

The first time you start Amarok in Fedora, the First-Run Wizard launches and displays the Welcome screen.

2. **Click Next to proceed to the Interface screen.**

Here, you can choose whether to have Amarok in one big piece or two different ones. I will assume you chose the former.

3. **Click Next again to take you to the Collection screen.**

4. **Here, expand (click the + symbols) and contract (click the - symbols) parts of the filesystem until you find the sections that contain your music.**

When you expose these folders (for example, `/home/dee/Music`), click in their boxes to add check marks. Leave the Scan Folders Recursively (look in subfolders), Watch Folders For Changes (pick up new files automatically), and Import Playlists boxes checked, or check them if they aren't. In this way, you let Amarok do most of the work.

5. **When you finish marking the necessary folders, click Next to continue to the Database screen. If you're not a database guru, just click Next and then Finish. Otherwise, take a moment to configure the database setup, and then click Next and Finish.**

Either way, when you finish, you reach the default Amarok screen.

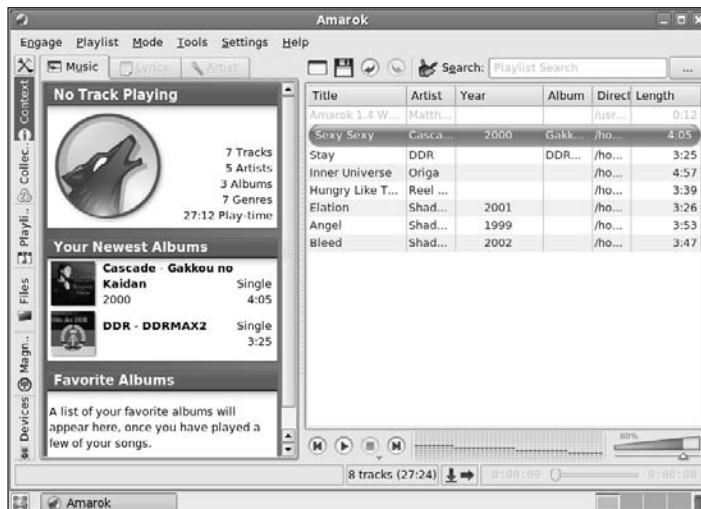
From here, you need to add music to the right-hand side in order to play it. You can individually drag and drop but there are lots of other ways to do this. Let me take you on a tour of the interface to point some of them out.

Along the left of the Amarok window is a collection of tabs. What you see and what I see might be different, depending on your version. Each tab offers a certain segment of functionality, and may change the interface to the left of the playlist. The tabs displayed in Fedora from top to bottom are:

- ✓ **Context:** Allows you to access information about your music (Figure 13-6). This tab actually has another three tabs inside it, along the top. The left-most tab, Music, shows you collated information about the music you have listened to up to this point. The Lyrics tab, in the middle, does its

best to load the lyrics for the song you're listening to and gives you clickable options if it can't decide exactly which song it might be — the first time you click this tab you may be offered a Run Script Manager button; if so, click it, select one of the two Lyrics options, and click Run. The rightmost tab, Artist, looks on Wikipedia's site (www.wikipedia.org) for information about the group or artist associated with the song. Lyrics and Artist both have controls along the top for you to explore.

Figure 13-6: The Amarok Context tab, Music subtab, with covers loaded by using the Cover Manager.



The Context tab is much cooler if you go to Tools → Cover Manager and click Fetch Missing Covers. When the proper cover is found, click Save. Or, if it wasn't found, click Next Cover to cycle through other possibilities until you find the right one.

- ✓ **Collection:** Shows your music collection. Above the listing, click the Group By button in order to choose features such as whether the collection is organized by group, album, and so on. In the Filter Here text box, you can type in keywords to determine what songs are displayed. Or, in the Entire collection dropdown-list box, you can specify if you want to have only the most newly added songs listed.
- ✓ **Playlists:** Here is where you create new playlists and access those that are already available. The Smart Playlists section offers you a variety of prebuilt playlists, such as All Collection for those who want to listen to everything they have, 50 Random Tracks to let the player pick them for you, Genres if you took the time to assign genre information when ripping your CDs, and more. Dynamic Playlists are different from Smart

Playlists in that they're constantly changing. The Random Mix dynamic playlist, for example, chooses music randomly from your collection and then keeps choosing it as you continue listening. For Suggested Songs, you put a few key songs in your playlist and then Last.fm (www.last.fm) builds you a personalized online radio station with music that it considers similar. If you're a big fan of online radio or are curious about it, then check out the Radio Streams section's Cool-Streams folder. Here you'll find lots of radio stations to choose from.

- ✓ **Media Device Browser:** Click this tab and then Connect to open the Configure Media Device dialog box. Here, you can choose from various portable storage device types if you have any (VFAT refers to Windows-based devices), so Amarok can load music from them.
- ✓ **Files:** Browse through your filesystem to find individual songs you want to click and drag into the playlist.

No matter how you get music over into the playlist on the right, when it's there you can click the Play button at the bottom of the pane to play the selected track. You can also double-click a track to play it. Below the volume slider you will find a set of controls (hover your mouse over the icon to see what it's currently set to), the most important of which are:

- ✓ **Random Mode:** Lets you turn playlist randomization on or off
- ✓ **Repeat:** Lets you set whether to stop playing music when the whole playlist has been played (no repeats), repeat a specific track, or repeat the whole playlist

There is much more to Amarok. Take the time to explore this tool and really customize it to your satisfaction. When you close this program, the Docking In System Tray dialog box appears. It warns you that Amarok will continue to run and appear as a system tray icon on the panel. You can avoid this by choosing Actions⇨Quit rather than clicking the X in the corner of the application to close the program. If you leave the program on the panel, you can open it at any time by clicking the little wolf icon in the panel.

Listening to Internet Radio

Listening to Internet radio is, for the most part, a snap in Linux, unless you want to listen to streams that are in Windows formats (see Chapter 15 for software you can purchase to solve this problem). The software options most people choose in general for Internet radio in Linux are RealPlayer, Amarok, and Mplayer (for Windows formats, for more about Mplayer see Chapter 14). Chapter 16 discusses how to install RealPlayer. The rest of this section

focuses on playing *streaming audio* (Internet Radio) in Amarok. For more on Amarok, see the previous section.



Mplayer is of questionable legality in North America and other places where intellectual property law and laws such as the Digital Millennium Copyright Act (DMCA) make it tricky for open source programs to navigate the licensing maze when it comes to music formats. For more on the DMCA and the problems it causes, see www.anti-dmca.org.

To find the preprogrammed stations in Amarok, select the Playlists tab and look under Radio Streams. Both the Cool-Streams and Shoutcast Streams sections have many selections to choose from. To add a new station to the list:

- 1. Click Add.**

A drop down list appears.

- 2. Select Radio Stream.**

The Add Radio Stream dialog box appears.

- 3. Enter the name for the stream in the Name text box.**

- 4. Enter the URL for the stream in the Url text box.**

- 5. Click OK to add the stream to Amarok.**

Your new stream appears at the bottom of the Radio Streams folder.

To listen to the station, double-click it in the listing. If you receive the message “No suitable input plugin,” then you are probably dealing with a Windows format.

Ripping Music Tracks from CDs

This is another topic that’s impossible to cover without at least acknowledging that both ethical and legal issues are involved. I’m not going to get into legalities here, but my personal ethics are that it’s fine to rip (copy) music off my own CDs for my own use. If I want to pull my favorite songs off CDs that I purchased and set them up so I can listen to them collectively in a random playlist off my computer’s hard drive, I don’t see a problem with this. However, doing this and then taking the CD back for a refund is theft, in my opinion.

So, with that said, to rip the tracks off of your CDs in Linux:

- 1. Choose Applications⇨Sound & Video⇨Sound Juicer CD Extractor (see Figure 13-7).**

Figure 13-7:
The Sound
Juicer CD
Extractor in
Fedora.




2. For each song that you don't want to rip, uncheck the check box next to the song.

The check mark disappears for each song that you don't want to digitize.

3. Choose **Edit** → **Preferences**.

The Preferences dialog box appears, as shown in Figure 13-8.

Figure 13-8:
The Sound
Juicer
Preferences
dialog box.



4. Select your preferred sound format in the Output Format drop-down list box.

If you're limited on space, I highly suggest Ogg Vorbis because this format is easily handled on Linux machines without any inherent legal problems and is also better compressed than MP3 for equivalent sound. However, if you want to go for higher quality for a home sound system, then chose FLAC, which is another open format but without any loss in sound quality.

5. Click the Folder drop-down list box to choose where you want to save your music files.

If the location you want to use isn't listed, select Other and then browse to the folder you want to use.

6. In the Track Names section, select how you want Sound Juicer to name and arrange the files in subdirectories.

7. Make any other configuration changes you want to make.

8. When finished making changes, click Close.

The dialog box closes.

9. Click Extract.

The Sound Juicer window shows you a progress bar along the bottom, and highlights the track it's currently working on. A dialog box appears when the extraction is complete.

10. Choose Eject in the dialog box.

The CD tray opens.

11. Remove the CD and close the tray.

12. Close Sound Juicer.

Burning Data CDs and DVDs

A CD and DVD burner is a great way to make data backups (especially if it's a rewriteable), save and share your digital photos, put together multimedia memento scrapbooks, and more. Many tools are available that allow you to burn CDs under Linux, but you can also burn CDs and DVDs by using the filesystem navigation tools discussed in Chapter 8, right-clicking the file you want to burn, and choosing the proper options from the context menu.

Here I will focus on K3b. When you launch this tool (using Applications→Sound & Video→K3b), you see the window shown in Figure 13-9.



Figure 13-9:
The K3b CD
and DVD
burner.

If all you want to do is create a CD or DVD out of an ISO file you downloaded, do the following:

1. Select either **Tools**⇨**Burn CD Image** if this is an ISO file for a CD-ROM, or **Tools**⇨**Burn DVD ISO Image** if this is a file for a DVD-ROM.

Both actions open the Burn CD Image dialog box (Figure 13-10).




Figure 13-10:
The K3b
Burn CD
Image
dialog box.

2. In the Image to Burn section, click the folder button and navigate to the file you want to burn onto CD or DVD.

3. When you find it, select it and click OK.

The filename and information now appears in the dialog box.

4. If you have more than one burner for some reason, make sure that the proper one appears in Burning Device, and change the selection if necessary.

5. In Speed, if you are burning a CD or DVD of a Linux distribution, lower it to at least 4x. Otherwise, you are usually safe to leave it as is.

If you find that when you burn CDs or DVDs they often seem to end up “coasters” (burns that didn’t work), lower the speed you usually use.

6. Click Save User Defaults if you want to use most of these settings for every burn.

7. Click Start to begin your burn.

A dialog box appears with two progress bars. The first shows you the write preparation progress, and then the second shows you the total progress, including burn time.

8. If this is a rewriteable media, then you’ll be asked if you want to erase its contents so it can be written over. Answer appropriately.

If you tell K3b to erase the media, then the Erasing dialog box appears during that part of the process. When the erasure finishes, the tray opens, ejecting the media, and then automatically closes again. This moment is when the actual burning session starts. The next time the disk ejects, the process is complete.

9. Remove your new CD or DVD from the drive and then click Close to close the burn information dialog box, and Close again to close the main burn dialog box.

10. Select File⇨Quit to close K3b.

Notice that this program has some other interesting features down in the bottom portion of the main window. Clicking one of the Project buttons will help you create a music (audio) CD, a data DVD, or a data CD from scratch.

Chapter 14

Messing with Video and Graphics

In This Chapter

- ▶ Watching online and downloaded videos
 - ▶ Watching DVDs
 - ▶ Creating and modifying high-powered graphics with The GIMP
 - ▶ Playing 3-D games
-

Video killed the radio star.

— *The Buggles*

There's no question that video is one area where Linux has been playing catch-up. The original Linux systems didn't even support a Windows-style desktop! In the early days of Linux, it was close to impossible to display any type of videos, let alone use them for playing movies or high-end computer graphics games.

Fortunately, nowadays, things are much different. There aren't many things your Linux computer can't do in the video world (unless of course it's illegal, but more on that later). In this chapter, I walk you through your Linux computer's video and graphics capabilities.

Watchin' Those Silly Web Clips

The Internet has become a clearinghouse for just about any type of video content you can imagine (and even some you couldn't imagine if you tried). From Google Videos to YouTube, there are lots of sites to keep you entertained.

Your Fedora system has the ability through the Firefox Web browser (see Chapter 10) to display lots of different types of Web clips, but unfortunately, not all of them by default.

Checking your plug-ins

The key to displaying Web video content in Firefox are plug-ins. As discussed in Chapter 10, plug-ins provide additional functionality to the Firefox browser.

You can easily check to see what plug-ins you already have installed in your Firefox setup by following these steps:

1. **Open Firefox by choosing Applications→Internet→Firefox Web Browser.**

You can also start Firefox by clicking the Globe icon on the top Panel.

2. **In the URL box, enter `about:plugins` and press Enter.**

The About Plug-ins page displays information about the current plug-ins installed, as shown in Figure 14-1.

By default, the Fedora 7 Firefox installation includes the following plug-ins:

- ✓ **Totem:** A full-featured video and audio player that can display AVI, ASF, WMV, OGG and MPG video formats, as well as play WAV and MP3 audio formats
- ✓ **Helix DNA:** A RealPlayer compatible player that plays RPM video formats
- ✓ **DivX Web Player:** For playing DivX-encoded AVI videos
- ✓ **QuickTime:** For playing Apple QuickTime MOV and QT formatted videos

Although on the surface it may seem that you've got most of your Web video bases covered, you may be disappointed with your actual video testing results. Although Totem supports the popular WMV and MOV video formats, there's more to quality video than just supporting the right file formats.

The Microsoft and QuickTime video formats use a proprietary encoding scheme called a *codec*. Because codecs are proprietary, they are strictly controlled under licensing agreements. Microsoft Corporation controls the WMV video format and, you guessed it, they don't provide a license for using this codec with most Linux systems. Because of this, Fedora doesn't support WMV video formats by default. The same applies to QuickTime MOV and QT video files.

That said, some resourceful open source programmers have reproduced the Microsoft video format codecs and created a software package called W32codec. Installing this package allows you to view almost every type of video available on the Web. Because this is not considered a legal software package in the United States, it is not included in the standard Fedora distribution. You can very likely find it by looking on the Internet, but don't say I didn't warn you.

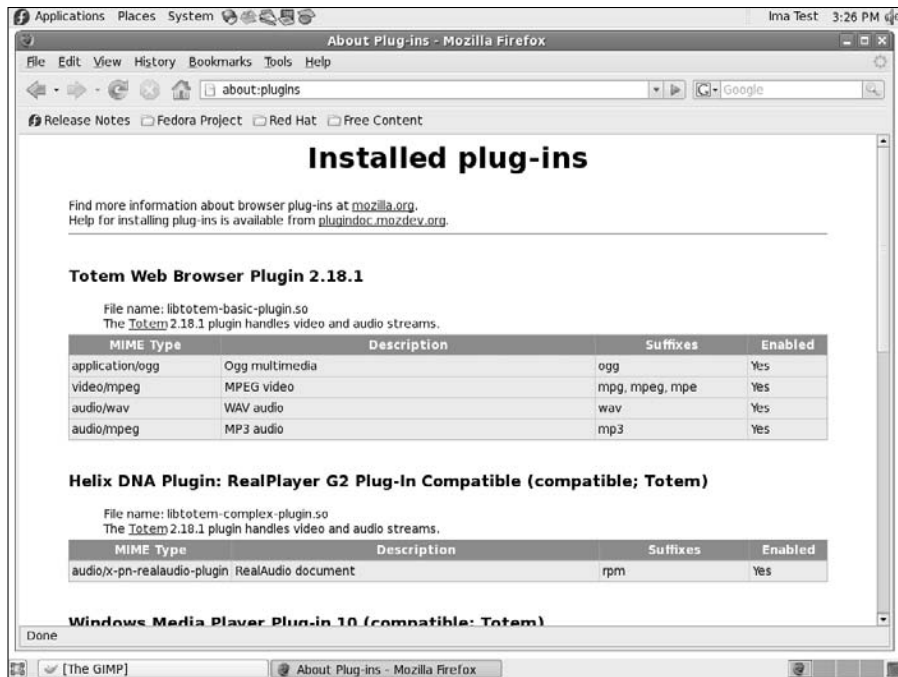


Figure 14-1:
The Firefox
About Plug-
ins page.



If you want to legally view Microsoft video formats in the United States in Fedora, you can purchase the CrossOver Office commercial software for Linux. This provides the appropriate codecs for viewing Macromedia Shockwave, Windows Media Player, and Apple QuickTime movie formats, as well as the libraries required to run many Microsoft applications directly on your Fedora computer. You can find CrossOver Office at www.codeweavers.com.

Adding plug-ins

A plethora of free video-related plug-ins are available for the Firefox browser. The clearinghouse for Firefox plug-ins is located at <http://plugindoc.mozdev.org> (the Firefox Web browser is based on the Mozilla open source browser). This site contains links to many different types of video-related plug-ins.

Many of the plug-ins listed in the Mozilla plug-in Web site are available as software packages in Yum. You can use the techniques presented in Chapter 16 to easily download new plug-ins for your Firefox browser.

The xine video player is a popular plug-in that supports lots of different video formats efficiently. Here's how to easily install it by using Yum:

1. Choose Applications⇨Add/Remove Software.

This starts the Package Manager tool. If you're not logged in as the root user, you are asked to enter the root user password.

2. Click the Search tab to search for a software package.

3. Type xine-plugin in the textbox and click the Search button.

4. Select the xine-plugin entry and click Apply to start the installation.

The xine-plugin package requires a few xine library files, which appear in the Dependencies list. Click Continue to install these library files as well.

5. After the installation finishes, open Firefox and view the about:plugins page.

You should now see that you have the xine plug-in available. Just look at all of those video types xine supports!

Getting even more video support

If you remember from Chapter 10, Firefox also supports a feature called *Extensions*. Extensions are small applications that provide additional features to your Firefox browser. If you're looking for more video support, or just some fancy video features, there are plenty of Extensions available to help out. The official Mozilla sites for Extensions is <https://addons.update.mozilla.org>. (Note that this is a site that uses the secure version of HTTP, so you must use https: rather than http:.)

One of my favorite Extensions is Unplug. Unplug allows you to save embedded streaming video content on a Web page, such as YouTube and MySpace, as a video file on your computer. Here's what you need to do to load and install it:

1. Open Firefox by choosing Applications⇨Internet⇨Firefox Web Browser.

You can also start Firefox by clicking the Globe icon on the top Panel.

2. Go to the <https://addons.update.mozilla.org> Web site and search for Unplug in all add-on pages.

At the time of this writing, it's located at <https://addons.update.mozilla.org/firefox/2254/>.

3. Click the Install Now link in the Unplug Web page.

Because the Unplug Extension is not signed, a dialog box will appear asking you if you really want to install it, as shown in Figure 14-2.

Figure 14-2:
Installing
the Unplug
Extension in
Firefox.



Click the Install Now button to install the Unplug Extension. When complete, the Add-ons dialog box appears, showing that the installation was successful.

4. Click the Restart Firefox button for the new Extension to work.

When Firefox starts, you should notice the cute green fish icon in your toolbar. To use Unplug, navigate to a Web page that contains an embedded streaming video, and then click the Unplug icon in the toolbar. Unplug scans the Web page looking for the video links, and then creates a separate Web page that contains what it thinks are the direct links to the video files, along with a button to save the video clip to your computer.

After you save the video clip on your computer, you'll probably want to be able to view it. That's where the Totem movie player comes in, which I talk about next.

Viewing Movie Files

With the popularity of digital video recorders, it's not uncommon to receive a video clip of your niece's ballet recital in an e-mail message, or from a file-sharing site. These files are often stored in a variety of video formats:

- ✓ **AVI:** Audio Visual Interleave format, a standard video format supported by most video software packages.
- ✓ **DivX:** A proprietary AVI video compression format that must be licensed to work in Linux.

- ✓ **MPG:** Moving Pictures Experts Group (mpeg) standard format. This includes both MPEG and MPG formats.
- ✓ **WMV:** Windows Media Video format, a proprietary Microsoft format.
- ✓ **RM:** The RealNetworks proprietary video format. This requires the RealOne video player.
- ✓ **MOV (and QT):** The proprietary Apple QuickTime video formats.

The default video file viewer in Fedora is Totem. You can start Totem either by double-clicking a saved video file, or launch it from the Panel menu by choosing Applications→Sound & Video→Movie Player.

As with the Firefox video plug-ins, due to licensing restrictions, the default Totem installation is somewhat limited in the types of files it can play. You can obtain the complete version of Totem, which supports more file types, by uninstalling Totem and installing the original software package. Just follow the installation steps in Chapter 16 to do this.

When you start Totem, the main viewing window appears, as shown in Figure 14-3.




Figure 14-3:
The main
Totem
viewing
window.

The Totem menu bar provides several options for viewing your videos. The **Movie** menu item provides options for loading movies into the Totem viewer:

- ✓ **Open:** Allows you to browse to a stored video file on your computer and load it into the Playlist
- ✓ **Open location:** Allows you to add a video file from a remote Web location to the Playlist
- ✓ **Properties:** Displays the properties of the video file, such as the format, and any embedded information such as the title, artist, and duration
- ✓ **Play/Pause:** Allows you to start a selected file in the Playlist
- ✓ **Quit:** Exits the Totem application

The **Edit** menu item controls how the movie is played:

- ✓ **Take screenshot:** Allows you to take a still snapshot of a scene in the video file
- ✓ **Repeat Mode:** Repeats the video file until you stop it
- ✓ **Shuffle Mode:** Plays video files in the Playlist in random order
- ✓ **Clear Playlist:** Removes all files from the Playlist
- ✓ **Preferences:** Allows you to set basic video and audio settings, such as setting color balance and audio output format

The **View** menu item controls how the video is displayed on the screen:

- ✓ **Fullscreen:** Displays the video file, using the entire computer screen.
- ✓ **Fit Window to Movie:** Automatically alters the Totem viewing window to fit the default size of the video file. You can also select to view the movie at 50% (1:2) of the default size.
- ✓ **Aspect Ratio:** Allows you to set the viewing mode to either standard 4:3 or 16:9 for widescreen videos.
- ✓ **Deinterlace:** Helps view some interlaced videos clearer.
- ✓ **Always on Top:** Keeps the Totem movie player the top window even if it's not the active window.
- ✓ **Show Controls:** Displays the Play, Skip, and Location controls at the bottom of the viewing window.
- ✓ **Subtitles:** Displays subtitles if included in the video.
- ✓ **Sidebar:** Includes the Playlist sidebar window in the viewing window.

The Go menu item controls which part of the video file is being played. For video files, it uses these options:

- ✓ **Skip to:** Locates a specific point in the video file
- ✓ **Skip Forward:** Goes quickly forward in the video file
- ✓ **Skip Backwards:** Goes quickly backwards in the video file

The Sound menu item controls how the audio is handled:

- ✓ **Languages:** Plays alternative language audio tracks if available in the video file
- ✓ **Volume Up:** Increases the volume level of the audio track
- ✓ **Volume Down:** Decreases the volume level of the audio track

The **Help** menu item displays the Totem manual, and information about the version.

Watchin' DVDs

Watching a DVD in Linux is a bit of a legal quagmire if you live in the United States. The Digital Millennium Copyright Act (DMCA) and other issues make it tricky for any open source program to navigate the licensing maze when it comes to movies that are encoded or protected in various fashions. However, not all DVDs have such countermeasures enabled: There are DVDs that Americans can watch under Linux with no trouble. (Note that I say *watch*, and not *copy* or *pirate*.)



For more on the DMCA and the problems it causes, see anti-dmca.org.

To watch a DVD in Fedora, you can use our friend Totem. However, as you saw when trying to view video files, the version of Totem that comes with Fedora is stripped-down in terms of formats it can support — again, this is a legal issue more than anything else. Often it is possible to replace this version with the full one by uninstalling Totem, adding software repositories that contain multimedia tools, and then installing the full version of Totem from the repositories (all these skills are discussed in Chapter 16). This is the solution I recommend for watching DVDs.



The Unofficial Fedora FAQ Web site (www.fedorafaq.org) provides some more tips on ways to view DVDs on your Fedora computer.

Creating and Modifying Graphics

The GIMP is a graphics program that's considered in many ways equivalent to Adobe Photoshop. Many don't consider *The GIMP* the friendliest program on the planet, but at the very least, it has enough features to keep you busy experimenting for weeks! To open *The GIMP*, choose Applications⇨Graphics⇨*The GIMP*.

When you start *The GIMP* for the first time, you have to walk through its user setup routine. Fortunately, you can just click Continue each time, unless you're a graphics guru who has a particular reason to want to do things in a different way. After you've clicked past all these dialog boxes, a collection of one or more dialog boxes pops up containing the *GIMP* main dialog box (see Figure 14-4) plus additional tool dialog boxes. I recommend closing all but the main one for now, just to keep your screen and brain uncluttered.

The GIMP is an incredibly complex program, with entire books written for the people who really want to use it heavily. Here are a few things to get you started.




Figure 14-4:
The GIMP's
main dialog
box in
Fedora 7.

GIMP basic tools

Because it's such an advanced tool, it's astonishing that the GIMP main menu bar only contains three items: **File**, **Xtns**, and **Help**. The beauty of GIMP is that all the tools appear in separate dialog boxes, allowing you to lay out multiple tools as you draw.

The File menu item allows you to open or acquire an image to work with. GIMP allows you to edit images from five different sources:

- ✓ Start a new image from scratch on a blank canvas.
- ✓ Load an existing image from your computer. GIMP supports all the standard image types, such as GIF, JPEG, TIFF, and others.
- ✓ Retrieve an image from a remote Web location.
- ✓ Retrieve an image from a scanner connected to your computer.
- ✓ Produce an image of a screenshot of your Desktop, or a single application window on the desktop.

When you open an image, it appears in its own editing window. If you open more than one image at a time, each image opens in its own editing window. This feature allows you to work on several images at the same time without having to flip around between images in a single window.

The top section of the GIMP main dialog box consists of the toolbox. The toolbox contains all the image editing tools you'd expect from an image editor. The toolbox consists of tools for selecting image areas; reorienting the image (such as flipping); drawing with a pencil, paintbrush, ink, or airbrush; filling colors, and adding basic special effects.

For each tool you select, the middle section of the GIMP main dialog box shows detailed settings for the tool. In Figure 14-4, the paintbrush tool is selected, so settings related to the paintbrush (such as brush size and opacity of the color) appear in the middle section.

You can also keep the detailed settings for multiple options open as separate dialog boxes. Choosing File⇨Dialog produces a list of the different tool dialog boxes you can have open on your desktop. This is shown in Figure 14-5.

Below the tools is an area to select the foreground and background colors. Double-click the black area to set the foreground color, and double-click the white area to set the background color for the image. Alongside that is an icon allowing you to easily change the drawing instrument shape and size.

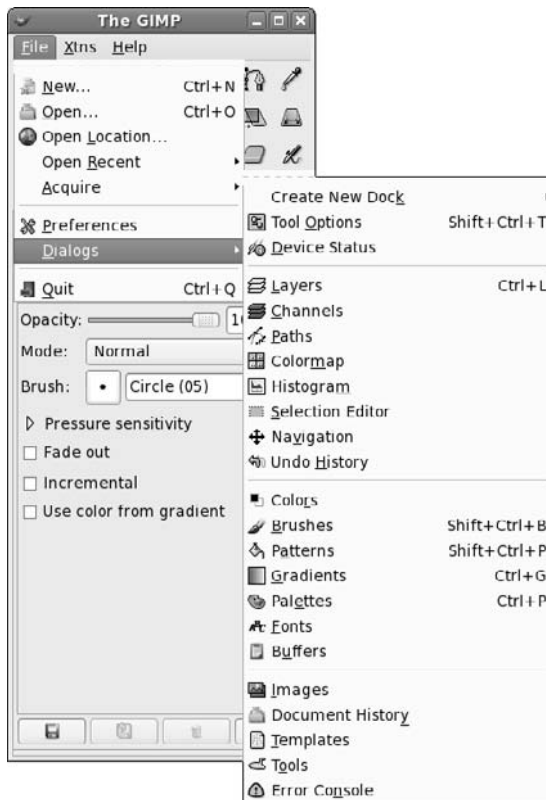


Figure 14-5:
The GIMP
tool dialog
box list.

To demonstrate the abilities of the GIMP, in the next section I do some work on a sample image.

Capturing screenshots


To get an image to work with, test out the screenshot feature. Follow these steps to create a screenshot image:

1. **Open GIMP by choosing Applications⇨Graphics⇨The GIMP.**
2. **From the GIMP menu bar, choose File⇨Acquire⇨Screen Shot.**

The Screen Shot dialog box appears, as shown in Figure 14-6.

You can select to capture either a single application window, or the entire desktop. You can also set a delay for when GIMP takes the screenshot, allowing you time to set up the window just the way you want it before you take its picture.

Figure 14-6:
The GIMP
Screen Shot
dialog box.



3. **When the screenshot is taken, GIMP opens the image in an image editing window. Select File⇨Save as to save the image file.**

GIMP allows you to save the image in most common image formats. All you need to do is place the proper image format extension on the file-name (such as .jpg, .tiff, .gif, or .bmp). After clicking OK, GIMP may produce another dialog box asking for options specific to the image file type you saved the image as.

You now have a saved image file that you can play with in GIMP.

Editing an image file

Now open your saved image file and play around with the GIMP image editing features.

1. **Select File⇨Open and select the image file you saved from your screenshot.**
The GIMP opens the image in an image editing window, as shown in Figure 14-7.
2. **In the GIMP toolbox, choose one of the selection tools (such as the rectangular box) and select an area on your image.**
3. **Again in the GIMP toolbox, choose the Rotate tool, and then click in the selected area on the image.**

The Rotate dialog box appears, allowing you to set the details of the rotation. As you increase or decrease the amount of rotation, you see the selected image area actually rotate in the editor. When you get to just the right spot, click the Rotate button to set the rotation.

4. **In the GIMP toolbox, select the drawing instrument of your choice (paintbrush, pencil, and so on) and do some doodling in the image.**




Figure 14-7:
The GIMP
image
editing
window.

By default, the drawing color is set to black. To change the color, you can double-click the foreground area in the GIMP toolbox. The Change Foreground Color dialog box appears, as shown in Figure 14-8.




Figure 14-8:
The GIMP
Change
Foreground
Color dialog
box.

This tool provides an easy way to select just the right color for your drawings. You can point-and-click to any shade of color you're looking for.

5. Save your new creation either as the same filename, or use another filename to create a separate image file.

That was some pretty basic stuff. You're probably thinking that you could do the same thing with some not-too-sophisticated tools, such as Microsoft Paintbrush. Well you're right, in a way. This does show how simple GIMP can be to manipulate images. But wait, there's more. Now follow along with this:

1. If you've closed your image file, open it by using **File**⇨**Open** so it's in a new image editing window.

The image editing window has a few tricks of its own that you can use.

2. Click the **Filters** menu item and select a special effect from the menu list.

The Filters section contains lots of special effects to enhance the look of your image. Figure 14-9 shows the result of the **Distorts**⇨**Whirl and Pinch** effect.

Figure 14-9:
The image
after
applying the
Whirl and
Pinch effect.



Now you're starting to see the power of GIMP. But wait, there's even more.

Using GIMP scripts

For those of us who are lacking in artistic ability, GIMP tries to help out as best it can. Lots of artistic people have contributed scripts to GIMP that provide more special effects than you'll see in a Hollywood action movie.

The key to special effects is being able to string together a series of special effects to create an overall effect. GIMP allows you to string special effects together by using a scripting language. Fedora includes both the Python-Fu and Script-Fu GIMP scripting languages.

The Python-Fu and Script-Fu scripts are accessible from the image editor menu bar. Each script has its own set of parameters that you can set to alter the effect of the script. Because scripts apply several layers of special effects to an image, they often take a while to complete.

The premade scripts included with GIMP produce some pretty amazing effects. Figure 14-10 demonstrates the Weave Script-Fu effect on the saved desktop image.




Figure 14-10:
The sample
image after
applying the
Weave
Script-Fu
effect.

That's pretty professional looking for a simple open source program! Take the time to play around with the different script effects available in GIMP. You'll be surprised at the quality of the special effects. As it says in the GIMP manual, most of the effects can be reversed by using the Edit→Undo menu item, so feel free to play.

Playing 3-D Games

One of the strengths of Windows PCs has been support for advanced video games. These games utilize high-end graphics that often require specialized 3-D video cards to operate.

Again, this is another area where Linux had never quite caught up to Windows, until now. Two major 3-D video card vendors, ATI and NVIDIA, have released Linux drivers for their products, allowing game developers to enter the Linux market.

There's a catch to this though. Both NVIDIA and ATI released Linux binary drivers for their 3-D video cards, but not the source code. A true open source system must include source code for all the components included in the system. This has caused quite a dilemma for Linux distributions.

For a Linux distribution to include the ATI or NVIDIA binary drivers violates the true spirit of open source software. But if a Linux distribution doesn't provide these drivers, it risks falling behind in the Linux distribution wars and losing market share.

Fedora has decided to solve this problem by splitting the difference. The Fedora 7 release doesn't include the ATI or NVIDIA 3-D graphics cards drivers. This makes the distribution completely open source. However, if you want support for your ATI or NVIDIA 3-D graphics cards, there's hope.

The <http://rpm.livna.org> Web site is a community-operated repository of nonapproved Fedora software packages. Not only does it include software packages for many of the video codecs I've been discussing in this chapter, but it also includes software packages for the ATI and NVIDIA 3-D graphics cards.

By default, the Fedora Yum software installer doesn't recognize the <http://rpm.livna.org> repository. The Web site includes instructions on how to add it to your repository list. After you do that, you're ready to install the ATI or NVIDIA 3-D graphics cards drivers:

- 1. Choose Application → Add/Remove Software.**

This starts the Package Manager tool. If you're not logged in as the root user, you'll be asked to enter the root user password.

- 2. Click the Search tab to search for a software package.**

- 3. Type the name of the desired video driver in the text box, and click the Search button.**

The ATI video driver is called `kmod-fglrx`; the NVIDIA video driver is called, strangely enough, `kmod-nvidia`.

- 4. Select the latest version of the video driver available from the list and click the Apply button.**

After you've installed the video driver, you need to activate it on the computer. To do that, you need to use your command line interface skills from Chapter 6:

1. **Start a command line prompt session by choosing Applications→System Tools→Terminal.**
2. **Become the root user by typing `su`, and entering the root password.**
3. **Shutdown the X server by using the `init` command.**

The command for this is

```
init 3
```

4. **Enable the 3-D video driver.**

To enable the NVIDIA driver, type:

```
nvidia-config-display enable
```

To enable the ATI video driver, type:

```
ati-fglrx-config-display enable
```

5. **Restart the X server by using the `init` command.**

The command to do this is

```
init 5
```

You should now have your 3-D graphics card working with Fedora! There are lots of 3-D games available in the standard Fedora software repositories:

- ✓ **freedom:** New episodes for the popular DOOM2 game
- ✓ **bzflag:** A tank versus tank battle game
- ✓ **openarena:** Based on the Quake III arena game
- ✓ **scorched3d:** Based on the Scorched Earth game
- ✓ **neverball:** A table tilt ball game
- ✓ **neverputt:** A miniature golf game
- ✓ **ppracer:** Racing Tux the penguin down a mountain slope

If you prefer some more exciting games, many popular gaming sites (such as icculus.org) provide Linux versions of lots of graphical open source games. Many of these games exploit advanced 3-D graphics to provide a realistic gaming environment.

If online 3-D games are your preference (or is that addiction?), you won't be left out in the cold with Fedora. Some of the most popular online 3-D gaming sites are starting to embrace the popularity of Linux and providing Linux clients.

Unfortunately, you'll need to do a little more work to get these clients installed on your Fedora computer. In this section, I cover the game Second Life as an example.

The Second Life online game is a virtual 3-D world entirely built by players. It uses the power of 3-D graphics cards to walk you through an amazing array of locations and lets you interact with graphical "people" that represent other players. After you enter the Second Life world, you can interact with others to purchase or trade products, buy land, and even visit entertainment establishments!

To enter the Second Life world, you'll need the client software. The software is available from the Second Life Web site as an installable binary package. You'll have to do some extra work to get the client installed, but it's worth it:

1. Open Firefox by choosing Applications→Internet→Firefox Web Browser.

You can also start Firefox by selecting the globe icon on the top Panel.

2. Go to the <http://www.secondlife.com> Web site.

3. At the bottom of the home page, click the Downloads link.

4. On the Downloads Web page, select the Linux link.

5. On the Linux downloads Web page, select the Linux i686 link.

The Fedora Download Manager starts. Select the save to Desktop option to save the client package for the software to your Desktop.

6. Exit Firefox.

7. Double-click the SecondLife installation package on the Desktop.

The File Roller application starts, showing the contents of the installation package.

8. In the File Roller window, click the Extract icon.

The Extract window appears. You can select where to extract the Second Life client folder to (most likely you'll want to put it under your Home folder).

After you have the software installed, it's time to play. Find the SecondLife folder you extracted from the installation package. In it is a file called secondlife. Double-click this file to start the client. The client connects to the Second Life site, and the game begins. Enjoy!

Chapter 15

Windows-Only Media Formats and Programs

In This Chapter

- ▶ Running Microsoft Office and Web tools with CrossOver Office
 - ▶ Using Windows inside Linux with Win4Lin
 - ▶ Playing Windows games with Cedega
 - ▶ Creating a “virtual computer” inside Linux with VMware
 - ▶ Trying out the free Windows inside Linux tool, Wine
-

A gentleman is a man who can play the accordion, but doesn't.

— Unknown

For us, nothing's more maddening than to find a file online or have someone send you a file, just to discover that you can't use it. Although this issue is less of a problem for Linux users as each day passes, things do sometimes still crop up that have to be worked around, particularly if you switch to Linux and then realize you have certain programs you can't live without and can't seem to replace.

Fortunately, you have a number of options available whether you need to use files or run Windows programs. If you're looking for software to try with existing files or a Linux equivalent to a beloved Windows program, a good place to start (aside from Chapter 16) is www.linuxrsp.ru/win-lin-soft/table-eng.html. Here, you can see why today it's surprising if you can't find a program under Linux to do what you need. Some of this software is “commercial” in that it's for sale, and other programs are free.

If you do find that you actually need to run a Windows program under Linux, both commercial and free options are available, although far more commercial options exist. This chapter looks at both possibilities so you can choose what meets your needs.

Commercial Software

A growing number of commercially available tools let you work with anything from Microsoft Office to your Windows games under Linux. Mind you, ultimately, the Linux user's goal is to not have to use any of these extra program layers — but in the meantime, it's nice not to be inconvenienced. I've reviewed, or at least used, every one of these products (often more than once over the years), so I'm not just going by their Web sites.



Many of these vendors offer server, as well as desktop, products. I'm focusing on the desktop. For more about their server products, check out their Web sites.

CodeWeavers

CodeWeavers (www.codeweavers.com) provides an excellent product that I use from time to time: CrossOver Office — don't let the name fool you, CrossOver supports games now as well. This program allows you to use Microsoft Office (along with a growing number of office applications) directly under Linux. It also comes with a hefty number of plug-ins (see Figure 15-1) that you might miss when it comes to Web surfing. It's a bit jarring to run Microsoft Word and Windows Media Player under Linux at first — because there's no *native* (without special software) way to do this otherwise — but it's certainly handy. You can even use iTunes through CrossOver Office!

What I most appreciate about CodeWeavers (and all these companies, really) is that they are very upfront about what works really well, what works except for a few features, and what doesn't work. All you have to do is check their Web site. For those who always seem to want to scream "piracy" when it comes to Linux, you actually do have to own the programs you want to run (say, Microsoft Office) and install them through CrossOver Office off their original CDs or DVDs. However, you don't have to own the version of Microsoft Windows the software runs under. CodeWeavers has built a reimplementa-tion of the necessary Windows functionality under the hood instead.

A plus to CodeWeavers CrossOver Office is that the applications run in individual windows just like every other program on your machine. A minus is that not all Windows software is supported because hooks have to be built under the hood for each new program.

Figure 15-1:
The
CrossOver
Office Install
Software
dialog box.



Win4Lin

Another option is Win4Lin (www.win4lin.com). Win4Lin actually offers you a platform within which you install Microsoft Windows (so have those CDs handy) and then run programs on top of it. At this writing, Win4Lin supports Windows 95, 98, ME, 2000, and XP (see Figure 15-2).

A plus for Win4Lin is that it supports a more full range of Windows software. By default, this program makes you run all of your Windows applications within a single window frame. See the Win4Lin Administration Guide for information on how to change this behavior.

Cedega

More interested in games than office tools? There's also TransGaming's Cedega (www.transgaming.com), which was once called WineX. This version of the free Wine project (see the upcoming section "Installing and Using Wine") is specifically designed to support Direct X, which is used in Windows to make programming for multimedia and games easier. Subscribers have access to regular updates, allowing them to make use of their favorite games and also add any games they like to the growing list of games that are being tested with the system.

Go to www.transgaming.com/searchgame.php to find a listing of the games, in order, that users have tried with Cedega (Figure 15-3) and what success they have had with these games.



Figure 15-2:
Windows
XP under
Win4Lin.

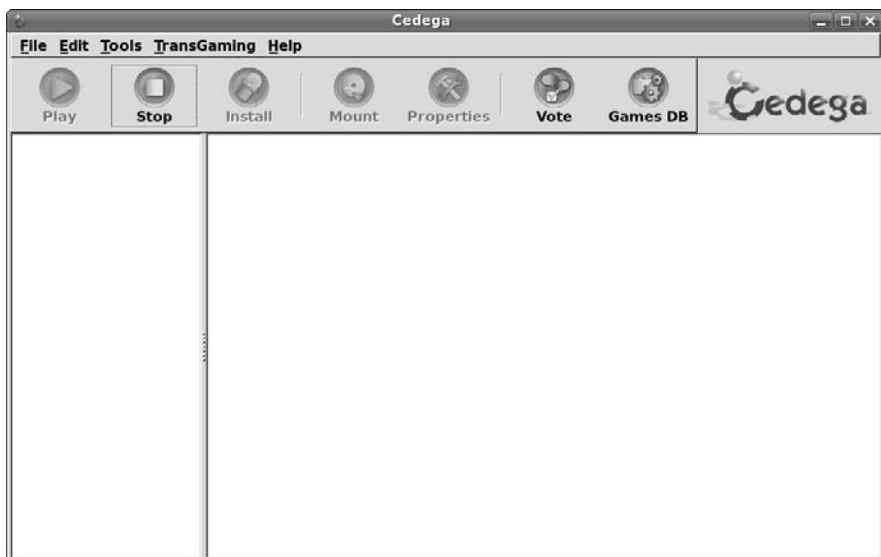


Figure 15-3:
Trans-
Gaming
Cedega.

A plus to Cedega is, again, that you don't need to keep those Windows (OS) CDs lying around. However, its support for individual games can be spotty, so be sure to check the Web site before you invest in this product. Also, be sure to make use of the help forums!

VMware

If you absolutely *have* to do something under Windows itself but only have one computer and it's running Linux (or vice versa, for that matter), there's also VMware (www.vmware.com). This is a popular tool in technical support offices where people need access to many different kinds of machines to test things. It's also popular with authors, like myself, because it allows us to grab those tricky screenshots of the boot process and installer screens.

In VMware, you run a full session of another operating system (or even an extra session of the same one you're running on the main machine) within your desktop — see Figure 15-4. If your machine is powerful enough, you could have five different operating systems all running in their own windows on a single computer.

A minus to VMware is that it requires a bit of extra oomph in your system. Check out its requirements and then try to go as far *above* them as you can. It also runs in a self-contained window. A plus, though, is that this program really is the next best thing to buying another computer. Much of the Fedora coverage within this book was gained by installing test/beta versions in VMWare.

There is a cool version of VMWare called VMWare Player (www.vmware.com/products/player/) that allows you to download a fully installed and set up system and play with it. You can test out a lot of Linux distributions this way, as a matter of fact. You'll find them in this section of the VMware Technology Network (www.vmware.com/vmtn/appliances/directory/cat/45).

Another free product is the VMware Server (www.vmware.com/products/server/). This tool allows you to install your own VMware sessions without all the fancy features that the pay product VMware Workstation has.



Figure 15-4:
A VMware
window
containing
Windows
XP, but
running in
Linux.

Installing and Using Wine

And then there's Wine (www.winehq.org), a complete rebuild of Microsoft Windows functionality that runs under Linux. Wine is in many ways the mother of everything in the previous section. However, because it's a free and very complicated project, it's not quite as friendly as the commercial versions. It's kind of like the others but with all the makeup and hairspray removed.

You can install Wine by using the information in Chapter 16. You will find that there are many packages that start with "wine." Install `wine`, `wine-cms`, `wine-core`, `wine-docs`, and `wine-tools` (or install all of them if you have lots of space and don't want to take any chances).



You can find the Wine documentation at www.winehq.org/site/docs/wineusr-guide/index. Start in the “Configuring Wine” section. When installed, you can find the Wine tools in Applications⇨Wine.

When you have Wine installed, you get to the fun part — trying to run your old Windows software. You don’t really run Wine, per se. Instead, you invoke Wine whenever you need to use something from the Windows world. To install a Windows program through Wine, insert the CD-ROM or floppy into the appropriate drive. Your distribution might open the File Manager for you, or you may need to open the File Manager manually. Regardless, your goal is to run the installation routine, which is typically `SETUP.EXE`. You can double-click this icon in the File Manager and, when asked what program you want to use to open this file, answer `wine`.


As I write this, I’m trying some old home-design program I bought back when I was silly enough to think I’d be able to afford to build a house before I was 90 years old. Because it’s difficult to describe a standard Wine session, I thought I’d walk you through how this goes for me.

The installer loads and runs just fine. I stick with the defaults (because Wine is still a work-in-progress and it’s best not to taunt it with getting too fancy) and then wait while the installer adds all the files to my SuSE system. Now, to run my new-old software. To run a Wine program, I first need to know what the path is to that program. I can find it by looking in the `~/.wine` directory. In this case, the program was installed into `/home/dee/.wine/fake_windows/Program Files/HOUSE`. The file `house.exe` looks like the one for launching the program.



If you’re trying to find the `.wine` directory in the GUI and can’t see it, you probably have to go to the View menu and choose Show Hidden Files. So, to run the program, I double-click it and, when asked what program to run it with, enter `wine`. That’s it (see Figure 15-5). Now pull out those old Windows 95 and 98 CDs and give them a try! Not all of them will work, but some probably will. The cool thing is that you can keep trying them on and off over time as newer versions of Wine come out.

Figure 15-5:
WINE
running an
old home-
design
Windows 98
program.



Chapter 16

Adding Software to Linux

In This Chapter

- ▶ Recognizing tarballs, RPMs, and compressed files
 - ▶ Creating tarballs and archives
 - ▶ Compressing files
 - ▶ Opening tarballs, archives, and compressed files
 - ▶ Installing and removing RPMs
-

I will make you shorter by the head.

— Queen Elizabeth I

When you start using a new operating system, one of the most frustrating things is trying to figure out all the goofy file extensions. The Windows world has `.exe` and `.zip`. The Macintosh world has `.bin` and `.hqx`. What about the Linux world? It certainly has its fair share of bizarre extensions, but, really, they make a great deal of sense after you know the programs that make them. In this chapter, you find out all about `.tar`, `.gz`, `.tar.gz`, `.tgz`, `.bz2`, and `.rpm`. Anyone up for a game of Scrabble with alphabet soup?

After you have the letter jumble all figured out, you'll be happy to find that Linux offers a number of cool tools for working with these crazy files, updating your system, adding new software, and more.

Opening Downloaded Files

The Linux and UNIX worlds are full of strange terms and acronyms. For example, if someone comes up to you out of the blue and starts talking about

tarballs, you probably get a mental image of sticky, smelly balls of tar, maybe rolled in feathers. Yet a tarball is something you run into regularly in the Linux world, especially when you're looking for software or you need to save yourself some space. A *tarball* is a bunch of files (and possibly directories) packaged together in a `tar` file and then compressed by using the `gzip` utility.

Fortunately, all you need to know is how to double-click a file in order to access the many formats listed in Table 16-1. When you double-click the file, your File Manager shows you what's inside.

Although Table 16-1 mentions operating systems, it doesn't contain hard and fast rules. People tend to use whatever kind of programs they're comfortable with no matter what operating system they're on.

Table 16-1 Potential Formats for Downloaded Files		
<i>Extension</i>	<i>Meaning</i>	<i>Program(s) Involved</i>
<code>.bz</code>	Older form of <code>.bz2</code> .	<code>bzip</code> , <code>bunzip</code>
<code>.bz2</code>	Slower but more efficient compression for some types of files, like text files.	<code>bzip2</code> , <code>bunzip2</code>
<code>.gz</code>	Typical compressed file for Linux and UNIX.	<code>gzip</code> , <code>gunzip</code>
<code>.iso</code>	A CD-ROM or DVD-ROM "image," which is a single file that contains a CD or DVD's entire contents. You have to tell your CD or DVD burner software that this file is an image so that it knows not to just place a copy of this file onto the media.	See Chapter 18.
<code>.rpm</code>	Not a file to "open"; see the "Installing New Software" section, later in this chapter.	<code>RPM</code>
<code>.tar</code>	A bunch of files bundled together.	<code>tar</code>
<code>.tar.bz2</code>	A <i>tarball</i> , which in this case is a <code>.tar</code> file inside a <code>.bz2</code> file.	<code>tar</code> , <code>bzip2</code> , <code>bunzip2</code>
<code>.tar.gz</code>	A traditional <i>tarball</i> , which is a <code>.tar</code> file inside a <code>.gz</code> file.	<code>tar</code> , <code>gunzip</code> , <code>gzip</code>
<code>.tgz</code>	A traditional <i>tarball</i> , which is a <code>.tar</code> file inside a <code>.gz</code> file.	<code>tar</code> , <code>gunzip</code> , <code>gzip</code>
<code>.Z</code>	Old-style UNIX compressed file.	<code>compress</code> , <code>uncompress</code>
<code>.zip</code>	Windows <code>.zip</code> file.	<code>zip</code> , <code>unzip</code>



WinZip (www.winzip.com) can handle `.gz`, `.tgz`, and `.tar.gz` files (along with the `.bz2` versions) for Windows users.

Compressing and Packaging Files to Share

Life isn't all about "take, take, take" (or at least I should hope not!). Sometimes you've just gotta give. Creating care packages to share with other folks involves finding how to tell Nautilus that you want to do so.

To package up and compress files for sending off to other people, navigate to the location where you've stored the file(s) — see Chapter 8 for how to move about in Nautilus — and then:

- 1. Determine whether you want to compress or package a single file, a group of files, or a whole folder.**

If the files and folders that you want to bundle together are flung all over the place in your filesystem, you may want to create a new folder and copy the items you want to bundle together into it — just for the sake of convenience. Chapter 8 shows you how to do so.

- 2. Select the item(s) you want to package.**

Chapter 8 explains how to do so. If you want to select a whole folder, navigate into its parent folder and just select the folder's icon instead of entering it.

- 3. Right-click over the item(s) (or choose the Edit menu).**

The context menu appears if you right-clicked or, in Nautilus, the Edit menu opens.

- 4. In the context or Edit menu, choose Create Archive.**

The Create Archive dialog box appears, suggesting a tarball (refer to Table 16-1) version of the file.

- 5. If you want a tarball, leave it alone, or change the file extensions (as discussed in Table 16-1) to match what you want.**

- 6. Click Create.**

Nautilus uses the end of the file's name to see what kind of compression to use.

7. **If you want to rename the file or change where it should go in your filesystem, do so now.**

Leave the `.tar.gz` or `.gz` part alone. Just change the first part. For example, if you're archiving the Files folder, the suggested name might be `Files.tar.gz`. If you want to use `Files100305.tar.gz` for the actual name, just add the date into the existing name.

8. **Click OK.**

The archive is created.

Updating Fedora

Fedora, by default, runs a utility that automatically reaches out to the Internet and checks for available operating system and software updates. You can either wait on this utility to tell you when something new is available, or you can update your system manually. Most people do a mix of the two, accepting automatic updates when it's convenient, and refusing them to go back later manually when it isn't convenient.

Automated updates

If you're connected to the Internet and an update becomes available, a white notification box appears in the upper-right section of your screen proudly proclaiming "Updates Available" and telling you how many. Click View Updates to open the Package Updater, and follow the instructions in the next section, "Manual updates," for how to update your computer.

Manual updates

If you closed the "Updates Available" box or just haven't tried to update your machine in quite a while, you can start the updater manually. To do so:

1. **Select Applications⇨System⇨Software Updater.**

This action opens the Package Updater (pup) dialog box, as shown in Figure 16-1 — you may first be asked to enter your root (administrative password). You will likely have to wait while the tool reaches out to the Internet and loads update information. When this process is complete, the progress dialog box closes and you can see all the available updates.



Figure 16-1:
The Fedora
Package
Updater.



If at this point you are told “Another application is running which is accessing software information,” wait a few minutes. The utility that tells you when updates are available may be in the middle of scanning.

2. **If you want information on any of the individual updates, click the item in the list and then click the arrow next to Update Details.**

Information about the update displays beneath the list of available updates.

3. **Click Apply Updates to proceed.**

The Resolving Dependencies for Updates dialog box appears with a progress bar — this process can take a while the first time you update the machine, as there may be one hundred updates or more.

When the tool is sure that everything that each program needs is taken care of, this dialog box closes and another may appear telling you that dependencies were added — *dependencies* are packages that another package has to have in place in order to install and run properly. If so, click Continue to proceed. If you see a dialog box saying that there were errors resolving the dependencies, or errors downloading *headers*, click Quit and try again later. The server might not be properly synchronized.

From here, you will see a Downloading dialog box. A progress bar shows you how much longer it will take for all the updates to be downloaded onto your computer. After downloading is complete, an Installing progress bar appears. When the updater finishes installing the packages, an Information dialog box appears, stating that the update is complete.

4. **Click OK to close the updater, if necessary.**

Installing New Software

In many cases, installing new software isn't much harder than updating, because Fedora has tools that let you install with just a few clicks. My focus in this section is on using these tools. Later, in the section "Finding More Software," I address how to add programs that aren't available through these tools. To add programs by using Fedora's Package Manager, make sure that you are connected to the Internet and then do the following:

1. Select Applications→Add/Remove Software.

You may be asked to enter your root. The Package Manager window appears (Figure 16-2). This window can take a minute or so to populate as the system looks to see what software it has installed and what is available.



Figure 16-2:
Fedora's
Package
Manager
with
Desktop
Environ-
ments
selected.

2. Determine how you want to choose your software:

- **Browse through categories:** By default you see the Browse tab. Software categories are displayed on the left, and to the right are groups of software that live within that category. An X in the box next to a group means that group is selected for installation, but it doesn't mean that *everything* in that group will be installed. Looking down and to the right will show you how many *packages*

(programs) from that group are actually selected. If you want to look at the packages and choose more, select the group you're interested in and click Optional Packages to see what else is available. Go through and click empty boxes to add Xs or click them again to remove Xs. When you're happy with your selections, click Close, and then move on to another category (on the left) or group (on the right). Then you can either select a new tab, or click Apply to install the software.

- **Search:** Click the Search tab (Figure 16-3). Select Available Packages so you can search through only the packages you haven't installed. Then type the text you want to look for in the text box, and click Search. The Search button changes to a Stop button, and as the Package Manager finds matches, they are displayed in the text area below, with the Stop button reverting to a Search button when the search is complete. If you feel you're ready to move on, click Apply to install your software.


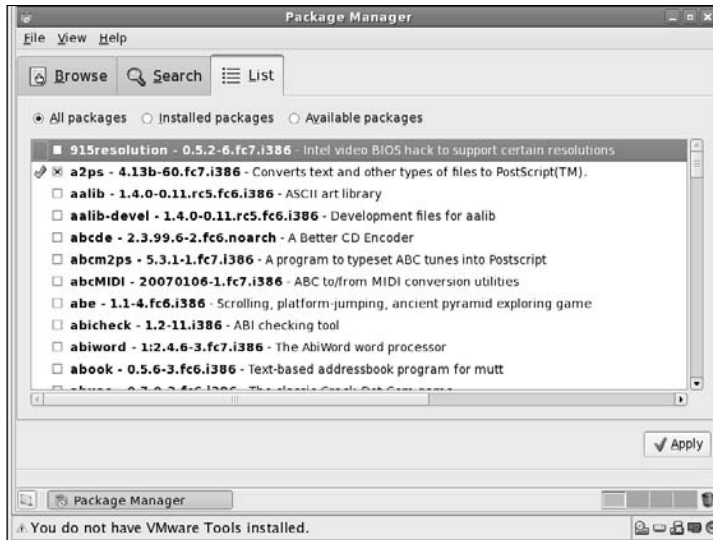


Figure 16-3:
Fedora's
Package
Manager
with the
Search tab
selected.

- **List packages:** Click the List tab (Figure 16-4), and after the list of all possible packages is loaded and displayed, click the Available Packages radio button. Now you can look through the list, clicking to add Xs to the boxes next to the software you want to install, or clicking again to remove Xs. If you now feel you have all the software you want selected, click Apply to install it.

Figure 16-4:
Fedora's
Package
Manager
with the
List tab
selected.



3. If you haven't yet, go through all three tabs if you want, mark with X's all the packages and groups you want to install, and then click **Apply**.

The Resolving Dependencies for Updates dialog box appears. This dialog box behaves exactly the same as the one in the Fedora updater, so you may be told that dependencies were added (click **Continue**) before the software downloads and installs. Eventually, the Information dialog box appears, telling you that the installation was completed.

4. Click **OK** to close the dialog box and package management tool.

Adding repositories

Fedora — along with other modern Linux distributions — uses *repositories* (sites that offer software to its automated tools) in order to track what's available and grab what it needs. By default, only the main Fedora and “updates” repositories are available for many reasons (some legal), so the first thing most people do is add a bunch of repositories to their list.

Wanting to spare you having to go through too much geeky (and potentially frustrating) typing, I scoured the Internet for an easier way. So, I present to you Max Kanat-Alexander's Fedora FAQ (www.fedorafaq.org). To add some of the most commonly used repositories, and find lots of cool tips and tricks, go to his site.

To add repositories from this site, go to the “Using Fedora and Installing Software” section, and then to “How do I install software in Fedora?” Click through for step-by-step instructions in the “Configuring Yum” section. Yum is the program that goes out to the Internet to check for software behind the scenes. Don’t worry about the other sections, such as “Using Yum,” unless you want to know how to get new software without using the tool discussed in the section “Installing New Software.”

Now that you’ve added more repositories, you may enjoy opening the Package Manager discussed in the previous section and seeing what new software is available.

Adding the software in this book

I discuss a number of programs in this book, and you may want to give them a try. Some are already installed by default. Others you need to add. You can add many through the Package Manager. For example, if you’re interested in trying the KDE desktop instead of GNOME, you can open the package manager, select Desktop Environments, and from there, select the KDE (K Desktop Environment) group. For individual programs:

- ✓ **Amarok:** Choose Applications⇄Sound and Video, click the Optional Packages button, and check the Amarok check box.
- ✓ **K3B:** Choose Applications⇄Sound and Video, click Optional Packages, and click the k3b check box.
- ✓ **Mplayer and Mplayer-plugin:** Add mplayer and mplayer-gui.
- ✓ **Thunderbird:** Choose Applications⇄Graphical Internet, click Optional Packages, and click the box next to thunderbird.
- ✓ **Java:** See the section “How do I install Java?” at www.fedorafaq.com.
- ✓ **Macromedia Flash:** Click the Search tab and search for flash-plugin.

Another program you may want to add is RealPlayer, which gives you MP3 support and can handle lots of different multimedia formats. To add this software:

1. **Point your Web browser to www.real.com/linux.**
You see Real’s Linux page.
2. **Click the Red Hat Package link beneath the big Download RealPlayer button.**
The Opening RealPlayer dialog box appears.
3. **Make sure that the Open with Install Packages option is selected.**

4. Click OK to download the file.

You're asked for your root password.

5. Type your root password and click OK.

The Completed System Preparation dialog box appears.

6. Click Continue to install the program.

When this process is complete, RealPlayer 10 appears in the Sound & Video menu.

Finding More Software

What if you can't find what you're looking for through the official (and not so official) sources discussed in the previous section? Those aren't your only options. Although I can't anticipate every situation you might find yourself in, I can at least give you some tips for how to find extra software and how to install much of it.

Here are the general steps for finding new software:

1. Check the Fedora FAQ (www.fedorafaq.org), Fedora News (www.fedoranews.org), the Fedora Forum (www.fedoraforum.org), and other online sites that people might recommend.

2. If you don't know what you want, find out by opening your favorite Web search engine and searching on a feature and the word `linux`.

For example, maybe you want something comparable to the program `irfanview` from the Windows world, so you would search on `irfanview linux`.

3. Sort through the search results and see whether a particular program is suggested. If not, add the word `equivalent` to your search and search again.

So, to continue the example, you would search again but this time use `irfanview linux equivalent`. Now you start to see a program called `xnview` mentioned. It wouldn't hurt to turn around and look and see whether Fedora's Package Manager offers this program, before you bother installing it by hand.

4. After you find a program you want, do a Web search for it.

More often than not, you find the program's home page.

5. Click through to that program's home page.

6. Click through the Download link on that page.

7. Locate and download the most specific version matching your distribution.

You may be offered, say, Windows, UNIX, and Linux options. You would choose Linux in that case. If offered Linux x86 versus Linux ppc, choose x86 unless you're using Linux on an Apple Macintosh computer (which is not covered in this book). If you're offered an RPM or a tarball (see the beginning of this chapter for more information on these), choose an RPM.

8. After you have the program downloaded, install it as follows:

- If it's an RPM, open your file manager and double-click the download in order to install it. You can also install it from a command line (see Chapter 6 for how to open a command line terminal). When you have the terminal window open, type `su` and press Enter. You'll be asked for your root password. Type it in, press Enter, and use the filesystem navigation commands in Chapter 7 to go to where you saved the RPM file, and then type `rpm -Uvh filename`. This command will install or update the program as needed.
- If it's a tarball, open your file manager. Double-click the file to open it up and look at its contents. There should be a file in there called README or INSTALL. This file contains instructions on what you need to do, and there may be more instructions available on the Web site itself. Working with tarballs just requires practice; it gets easier over time, so extract the file and get to it!

Upgrading Your OS

Typically, Fedora users upgrade by downloading the new version of the distribution (say, Fedora 8), starting the installation, and selecting Upgrade rather than a new install.

Chapter 17

Basic System Administration

In This Chapter

- ▶ Managing users
 - ▶ Playing with printers
 - ▶ Watching programs
-

After one look at this planet any visitor from outer space would say “I want to see the manager.”

— William S. Burroughs

Okay, so you’ve got your Fedora Linux system running exactly the way you want it, with just the right combination of software packages installed. You’re writing documents, creating fancy graphics, listening to music CDs, and even playing games. But don’t sit back and relax: There’s still more left to do.

You’ll always need to be tweaking things on your system — whether it’s allowing other people to use it, installing a new printer, or trying to figure out why a software program is running so slow (yes, even Linux programs can go bad sometimes). Keeping everything in order and running properly can be a full-time job. Even for a home Linux system, you may need to manage several user accounts and printers, as well as keep track of over a dozen programs running on the system.

Fedora provides a few tools for us that help in the Linux management process. No doubt you’ve seen the Administration section under the System menu area on the desktop (see Chapter 5). Some of the items are fairly self-explanatory, such as Date and Time settings, Display settings, Keyboard settings, and the Language setting. However, other items aren’t quite as easy to figure out how to use. This chapter walks through a few of the administration tools Fedora provides to help make running your system a little easier.

Users and Groups

Selecting System⇨Administration⇨Users and Groups from the desktop starts the Fedora User Manager tool. If you aren't already logged in as the root user, you are asked to enter the root user's password before it starts. The main User Manager window appears, as shown in Figure 17-1.

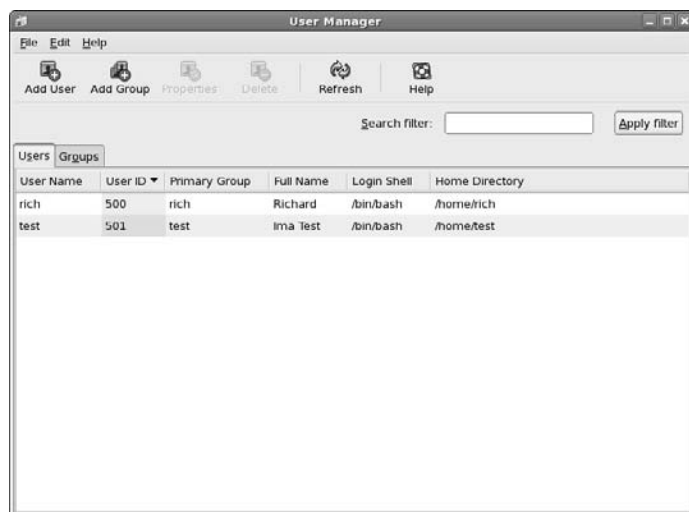


Figure 17-1:
The Fedora
User
Manager
tool.

Whether you support a Linux system with hundreds of users, or just have a user account for you and one for your cat, you'll need to know how to manage user accounts. This is done with the Fedora User Manager tool.

The main window contains two tabbed panes — one for managing users, and the other for managing user groups, plus several helpful buttons in the toolbar.

The menu bar provides text access to the same features available on the toolbar, plus one additional feature not available on the toolbar. Selecting Edit⇨Preferences from the menu bar produces the Preferences dialog box, as shown in Figure 17-2.




Figure 17-2:
The User
Manager
Preferences
dialog box.

This dialog box contains settings that control the behavior of the User Manager tool itself:

- ✓ **Hide System Users and Groups:** The Fedora system requires lots of additional usernames to control background system programs. Messing with these users and groups isn't recommended, so hiding them helps clean up the user viewing area, and helps prevent any accidents from happening (like deleting an important system user account).
- ✓ **Automatically Assigned UID Must Be the Highest:** Fedora assigns each user account a unique User ID number (UID). The User Manager is more than happy to automatically assign UIDs to new users, or you can choose to manually assign your own UIDs. This setting ensures that if the User Manager assigns UIDs, the numbers are in order after the UIDs for any accounts you've already created. It's safe to keep this value set unless you're doing your own UID assigning.
- ✓ **Automatically Assigned GID Must Be the Highest:** Fedora also assigns each group a unique Group ID number (GID). This setting ensures that any automatically assigned group numbers are in order at the end of the list. Again, it's okay to keep this setting enabled.
- ✓ **Prefer that Private Group GID Is the Same as UID:** This feature ensures that each user account's UID is the same number as the assigned group ID. This helps simplify user management by ensuring that each user account belongs to its own group with the same number.

After you set the User Manager tool preferences, you're ready to actually start managing users and groups.

Adding new users

When you first installed Fedora, it created a user account for you. That account has its own home folder for storing files. If you are in an environment where others will use the same system, it's a good idea to create separate user accounts for each user. That way, you don't have to worry about the cat accidentally deleting your important documents.

If you need to add additional user accounts, here's what you need to do:

- 1. Click the Add User button.**

The first button in the toolbar allows you to add a new user to the Fedora system. Clicking that button produces the Create New User dialog box, as shown in Figure 17-3.

- 2. Type a username for the user in the User Name text box. Type the full name of the user for documentation purposes in the Full Name text box.**

Figure 17-3:
The Create
New User
dialog box in
the User
Manager
tool.



3. **Type a password for the user in the Password text box, and then retype it in the Confirm Password text box.**

For security you should always use passwords that aren't common words, contain a mixture of numbers and upper and lowercase letters, and are a minimum of eight characters long.

4. **Select a login shell for the user.**

Fedora Linux supports several different login shells. The login shell allows you to work with text commands and scripts (see Chapter 6 for help with working with the login shell). The default login shell in Linux is the Bourne again shell, called `bash`. Using this default is just fine.

5. **Create a home directory for the user.**

By default, the User Manager tool creates a directory for the new user to store personal files, located under the `/home` folder, and named by using the user account name. You can elect to not create a home directory, or to create the home directory in an alternative location. Unless you're running a huge data center that relocates home folders to another disk drive, the suggested default location is just fine.

6. **Create a private group for the user.**

Each user must belong to a default group. Some Linux distributions create a single group that contains all users. By default, Fedora creates a unique group for each user, ensuring that users can't accidentally share files. It's fine to allow Fedora to do that.

7. Specify user ID manually.

By default, the User Manager tool automatically assigns a UID for the new user account. If you prefer to use your own numbering system for UIDs, you can assign your own UID here.

8. Click the OK button to create the new user.

After the new user account is created, it appears in the user list under the User tab.

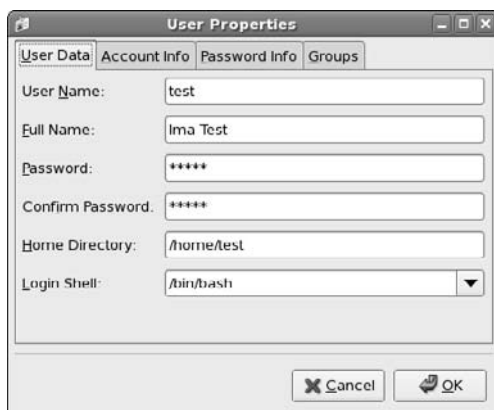
Modifying existing users

A few user account settings are not assigned when you create a new account. You can get to those from the User Properties dialog box:

1. Either double-click a user account listed in the User tabbed pane, or select the user account and click the Properties button on the toolbar.

The User Properties dialog box appears, as shown in Figure 17-4.

Figure 17-4:
The User
Properties
dialog box in
the User
Manager
tool.

**2. Alter data on the User Data tab.**

The User Data tab contains the same fields as shown in the Create User dialog box. You can alter the username, full name, home directory, or login shell here, as well as change the password for the user account.

3. Click the Account Info tab to set an account expiration date or lock the account.

The expiration date doesn't mean the user has expired, just the login account has. The account expiration date locks (but doesn't delete) the account when the date is reached, and the user is unable to log in. If you suspect unauthorized activity from this user account, you can also manually lock the account, preventing the user from logging in, but keeping the account still configured on the system.

4. Click the Password Info tab to force password expiration rules.

This section allows you to set whether users must change their passwords. There are fields for how many days users have to change their passwords, how many days before the password expires that the user is warned, and how many days after the password change date the account is locked if the password is not changed.

5. Click the Groups tab to assign the user account to additional groups.

By default, when a user account is created Fedora automatically creates a unique group for the user. Here's where you can assign the user account to other groups as well. All the groups defined on the system appear within the tab. Select each group the user account should belong to. You can also change the primary group the user account belongs to (by default, it's the same name as the user account).

6. Click OK to accept any changes made to the user account.

Any changes made to the user account take effect immediately. However, if you've added a user to a new group, permissions often don't work properly until the user has logged out and then back in again.

Deleting a user

Although it's quite an ordeal to add a new user account, removing user accounts is a snap. Just select the user account you want to remove from the User tab list, and then click the Delete button in the toolbar. A simple message appears asking you if you're sure of your selection. After double-checking to make sure you're deleting the correct account, just click Yes.

That's all there is to user management. In the next section, we tackle group management.

Adding new groups

Groups help you manage access to files and folders on the system. Each file and folder contains three sets of permissions — one set for the owner, another set for the group the file or folder is assigned to, and the last set for

all other users on the system. By assigning a file or folder to a special group, you can allow all users that belong to that group to have full access to the file or folder.

Here's what you'll need to do to create a group in User Manager:

- 1. Click the Add Group icon in the toolbar.**

This starts the Create New Group dialog box, as shown in Figure 17-5.

Figure 17-5:
The Create
New Group
dialog box
in the User
Manager
tool.



- 2. Type the name of the new group in the Group Name text box.**

- 3. Select if you want to manually specify the Group ID.**

By default, User Manager automatically assigns a unique group ID number (GID) to the new group. If you prefer to use your own GIDs, select the check box, and then select an available GID from the drop-down list.

- 4. Click OK to accept the settings and create the group.**

The new group is created, and appears in the group listing under the Groups tab on the main User Manager window.

Modifying groups

Now that you've created a new group, most likely you'll want to add some users to the group. (What good is a group with no members?) Do this by using the Properties dialog box:

- 1. Double-click the group name from the Group tab listing, or select the group name from the listing and click the Properties toolbar icon.**

The Group Properties dialog box appears, as shown in Figure 17-6.

- 2. Click the Group Data tab to rename the group.**

If you want to rename an existing group while keeping the same GID, enter the new name in the Group Name text box.

Figure 17-6:
The Group
Properties
dialog box in
the User
Manager
tool.



3. Click the Group Users tab to add users to the group.

The Group Users tab contains a drop-down list box containing all the configured users on the system (even hidden system users). Select the check box for each user you want to add to the group.

4. Click OK to apply the changes to the group.

You should see any newly added user accounts appear on the User Management Groups tab under the group name.


Printing

Selecting System⇨Administration⇨Printing from the desktop menu starts the Fedora Printer Configuration tool. Even though technology gurus keep predicting a paperless society, for now, we're still stuck having to print some things out. In the past, printing was one of the dark areas in Linux. Trying to get modern printers working with Linux was quite the challenge. However, some amazing advances have recently been made in making Linux more printer friendly.

Possibly the biggest advance in this area has been the Common Unix Printing System (CUPS). CUPS provides a common interface between UNIX (and Linux) systems and printers. It runs in background as a service, connecting to any defined printers, and waiting for applications to send print jobs. Because it runs in the background, it can both talk with remote printers, and accept print jobs from remote systems.

The Printer Configuration tool provides an easy way to configure both the CUPS server running on the system and any printers you've defined. The main Printer Configuration window is shown in Figure 17-7.

Figure 17-7:
The Fedora
Printer Con-
figuration
window.



There are two sections to the Printer Configuration window:

- ✓ **Server Settings:** Allows you to control how the CUPS server manages system printing features
- ✓ **Local Printers:** Lists all the locally defined printers on the system, and allows you to change individual printer properties on each of them

Selecting each section produces a different set of properties to configure in the right-side window area. There are also six buttons available in the toolbar:

- ✓ **New Printer:** Starts a simple wizard to guide you through adding a new printer to the system.
- ✓ **New Class:** Starts another simple wizard that guides you through grouping configured printers together to form a printer pool (called a *print class* in CUPS). Sending print jobs to a print class enables the job to be printed by any available printer in the class.
- ✓ **Copy:** Copies an existing local printer definition to create another printer of the same make and model.
- ✓ **Delete:** Deletes an existing printer.
- ✓ **Goto Server:** Enables you to connect to a CUPS server running on a remote computer.
- ✓ **Refresh:** Refreshes the printer display area, showing any newly discovered or added printers.

The following sections describe how to use the Printer Configuration window to set the CUPS and printer properties for your system.

Printer server settings

Clicking the Server Settings item on the Printer Configuration window produces the Basic Server Settings window — refer to Figure 17-7. There are a few different settings you can play with here to help out with your printer administration:

- ✓ **Show Printers Shared by Other Systems:** Displays printers found by browsing the network.
- ✓ **Share Published Printers Connected to This System:** Allows remote clients to connect to any of the local printers marked as shared (more on that later).
- ✓ **Allow Remote Administration:** Enables remote clients to connect to the CUPS server running on this system.
- ✓ **Allow Users to Cancel Any Job (Not Just Their Own):** By default, normal users can only cancel their own print jobs. Enabling this feature allows any user to cancel any other user's print job. Although this is a handy feature, it can get dangerous in a huge multiuser environment (especially if your users like playing tricks on one another).
- ✓ **Save Debugging Information for Troubleshooting:** By default, the CUPS server generates a moderate amount of logging information to monitor usage or problems. If you are having trouble with a specific printer configuration, you can enable this feature to produce more (lots more) information in the log files.

After determining the settings appropriate for your CUPS server environment, you can move on to creating the individual local printers.

Adding a new printer

One amazing feature of Fedora is its ability to automatically detect printers connected via USB cables. If you have one of these printers, you most likely don't need to add it to the system. You should already see it listed under the Local Printers listing. You can just move on to the next section to configure it.

If you are not fortunate enough to have a USB-connected printer, you'll have to manually add your printer. Here are the steps for doing that:

1. Click the New Printer button on the toolbar.

The New Printer Wizard appears, as shown in Figure 17-8.

Create a unique name for the printer. You may also assign a description and a location to the printer to help users know exactly which printer it is (this is usually only helpful in large organizations that have lots of printers). Click the Forward button to move to the next wizard window.

Figure 17-8:
The New
Printer
Wizard in
the Printer
Configura-
tion tool.

New Printer

Printer Name
May contain any printable characters except "/", "#", and space
myprinter

Description (optional)
Human-readable description such as "HP Laserjet with Duplexer"
A test printer

Location (optional)
Human-readable location such as "Lab 1"

Cancel Forward

2. Select the connection type of the new printer.

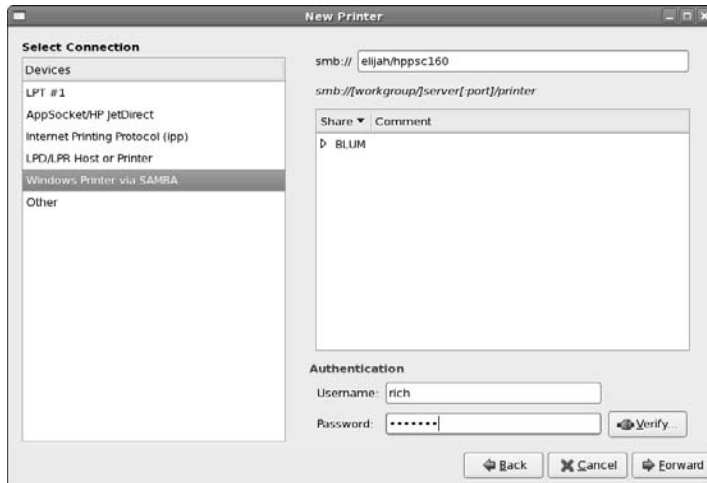
Although the printer is defined as a local printer, this utility also allows you to configure six different types of printer connections to make a printer available to the system. These printer connection types are listed in Table 17-1.

Table 17-1 **Fedora Printer Connection Types**

<i>Connection Type</i>	<i>Description</i>
LPT#1	A printer connected via the 25-pin LPT port on the local computer
AppSocket/HP JetDirect	A printer directly connected to the network, using a network card and Internet software
Internet Printing Protocol (IPP)	A remote CUPS system advertising local printers that can be used remotely
LPD/LPR Host or Printer	An older UNIX standard for remotely sharing local printers
Windows Printer via Samba	A Microsoft Windows workstation or server that advertises shared local printers
Other	A specific Universal Resource Identifier (URI) used to map to a network printer

As you select each connection type in the left side of the window, a different Properties area appears in the right side. For example, if you select Windows Printer via Samba, text boxes appear where you can enter the printer name, and a user ID and password required to access the network printer, as shown in Figure 17-9.

Figure 17-9:
Using the
New Printer
Wizard to
map to a
Microsoft
Windows
shared
printer.



After you've selected the connection type and changed any Property settings, click Forward to continue on with the wizard.

3. Select the printer manufacturer, or the location of the PPD file.

Fedora uses standard Postscript Printer Description (PPD) files to handle formatting print for printers. These use the same concept as the standard printer drivers for Microsoft Windows you're probably used to. Each printer must have a PPD installed for CUPS to know how to format text and graphics sent to the printer.

This wizard window allows you to select the PPD file to use for the new printer. You have two choices:

- Select the printer manufacturer from the list of installed drivers
- Select to install your own PPD file for the printer

If you are lucky enough to have the PPD file for your printer, copy it to a location on your Fedora system and select the Provide PPD file option. Browse to the location of the file and select it, and then click Forward.

If you don't have the PPD file for your printer, you have to hope that Fedora has your specific printer make and model defined in its library. If you find the printer manufacturer listed, select it and click Forward.

4. If you selected a printer manufacturer, the next wizard window provides a list of specific printer models and PPD files. Select the printer model and (optionally) the proper PPD file.

This wizard window asks you to select the specific model type for your printer. Hopefully your printer model is listed. If not, you must go back a step and find your own PPD file to install.

If your specific model is listed, select it, and a list of available PPD files is shown. Some models may only have one PPD file, but others may have two or more files to choose from. In that case, one is usually marked as recommended. Try that first: If it doesn't work, select a different PPD file. Clicking the Forward button takes you to the final wizard window.

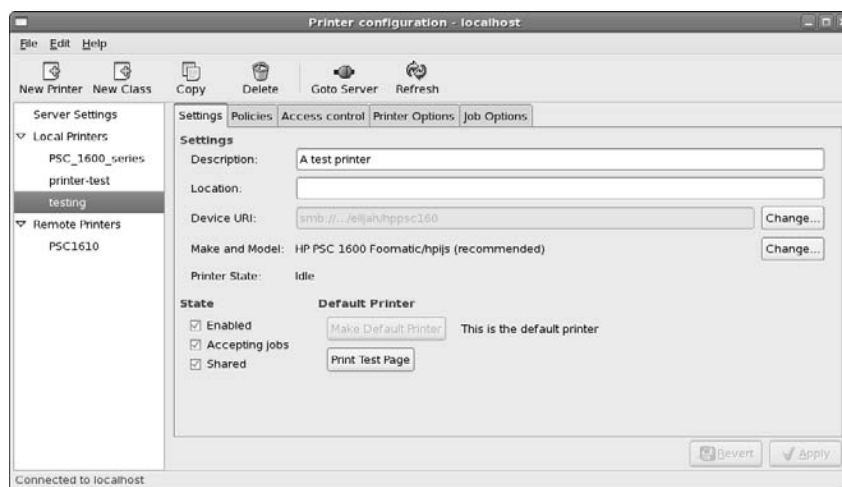
5. Click Apply on the summary window to create the new printer.

The new printer is added to the list of local printers (even if you mapped to a remote printer). You should now see the new printer when printing from applications on your system. Before getting too carried away with your printing, it's a good idea to check how the printer is configured.

Modifying printer properties

You can modify the properties of any printer on the system (whether Fedora automatically created it or you manually created it). Select the printer you want to configure in the left window, and the printer properties appear in the right side, as shown in Figure 17-10.

Figure 17-10:
Modifying
printer
properties in
the Printer
Configura-
tion tool.



There are five different tabs that contain printer information you can modify. Those tabs are Settings, Policies, Access Control, Printer Options, and Job Options. In the next few sections, we take a closer look at each.

Settings

The Settings tab provides some basic settings you can change for the printer. From here, you can change the description and location tags for the printer, the URI of the printer, and the PPD file used for the printer.

There are three check boxes that enable you to control the state of the printer:

- ✓ **Enabled:** Allows you to take the printer offline if there's a problem
- ✓ **Accepting Jobs:** Allows you to temporarily suspend processing print jobs
- ✓ **Shared:** Allows you to set whether this printer is a shared resource on the network

Besides the three check boxes are two other buttons:

- ✓ **Make Default Printer:** Makes this printer appear first when printing from applications
- ✓ **Print Test Page:** Provides an easy way for you to test if the printer and PPD file are working properly

These cover the general settings for the printer. The other tabs dive into more advanced settings.

Policies

The Policies tab contains four policies (or rules) that control printer behavior. Two of the policies control banner pages for print jobs, and the other two control how the printer operates.

Banner pages allow you to print a special page describing the print job. The Starting Banner — a special page printing out before the print job — allows you to print a cover sheet for the print job. Theoretically it's supposed to provide some privacy, blocking people from seeing the first print page on the printer, but really, who wouldn't just look under the banner page?

The Ending Banner page allows you to print an ending page marking the end of the print job. If you're in a high print-volume environment where lots of print jobs are pouring out of the printer, having a starting or ending banner helps keep everyone's print jobs separate.

There are two operation policies that you can set:

- ✓ **Error Policy:** Determines how the printer reacts to an error in the printing process. The choices are Abort the Print Job (giving up all hope of printing the job), Stop the Print Job (possibly trying again after manual intervention) or Stop the Printer (preventing all other print jobs from printing). The default value is to stop the printer, on the theory that if one print job errors out, any others that follow are likely to do the same.
- ✓ **Operation Policy:** Determines the mode the printer runs in. At this time, the only setting for this is default.

Access Control

The Access Control tab provides a method for you to restrict access to the printer. There are two ways to do this.

First, you can list every user account that the server will prevent from using the printer. This assumes that any user account not in the list can print, and any user account on the list can't print.

The second method is to list the user accounts that are allowed to use the printer. This assumes that any user account not in the list can't print. Obviously, which one you choose depends on if you have more people you want to allow to print or to restrict from printing.

Printer Options


The Printer options tab allows you to set some properties for the physical printer setup. The properties available for you to modify are based on information provided by the individual printer's PPD file and depend on the physical characteristics of the printer.

These properties are divided into separate categories, depending on the printer's capabilities. The options available for my test printer are shown in Figure 17-11.

The General settings handle properties such as what types of paper the printer can handle, the print qualities it can produce, and the number and types of input trays. You can force the printer server to request a specific paper size for all print jobs, or even request which paper tray to take paper from.

The Printout Mode settings set the default print quality used by the printer. Again, this setting depends on the capabilities of the particular printer, but usually there's a range of Dots Per Inch (DPI) settings and color settings for you to choose from.

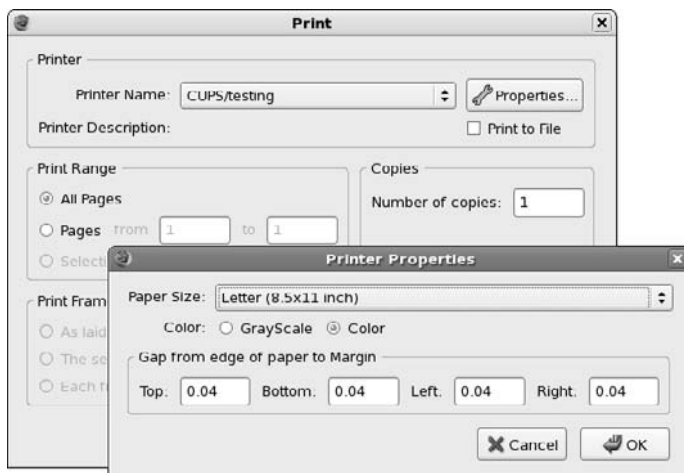
Figure 17-11:
The Printer
Options tab
of the
Printer
Configura-
tion tool.



Job Options

Finally, the Job Options tab allows you to set the default properties for print jobs sent to the printer. When you request an application to send something to the printer, a standard GNOME Print Dialog box appears, as shown in Figure 17-12.

Figure 17-12:
The
standard
GNOME
Print and
Properties
dialog
boxes.



You can select several printing properties for the specific print job, such as whether to print in landscape mode, how many copies of the document to print, the paper size to use, and so on. If you prefer to use any of these settings by default, you can set that in the Job Options tab.

Setting the Job Options is somewhat odd:

1. On the Job Options tab, select the job option from the Specify Default Job Options drop-down list box at the top of the window.
2. Click the Add button to add it to the list of available options to modify.
3. Change the default values for the option.

The completed list of options appears within the Printer Configuration properties window, as shown in Figure 17-13.

Figure 17-13:
The Job
Options tab
in the
Printer
Configura-
tion tool.



If you decide to revert to the default setting for an option, click the Remove button. The option is removed from the options list.

System Monitor

Selecting System → Administration → System Monitor starts the System Monitor tool. Although things usually go well on the Fedora system, sometimes applications can go awry. When that happens, it's handy to be able to watch what's going on under the hood of the operating system.

The System Monitor tool provides information about how the system hardware and software are working. There are four window tabs in the System Monitor:

- ✓ **System:** Shows the host name, hardware information, and system status (available hard disk space)
- ✓ **Processes:** Shows information about programs currently running on the system
- ✓ **Resources:** Shows information about CPU and memory usage
- ✓ **Filesystem:** Shows information about what hard drives are mounted on the system

Of course, just having a bunch of information thrown at you doesn't help at all. You need to know what that information means for it to be meaningful. The following sections help you understand how to interpret the various pieces of information contained in the System Monitor.

Processes

Clicking the Processes tab provides an overview of the programs currently running on the system, as shown in Figure 17-14.




Figure 17-14:
The
Processes
tab in the
System
Monitor
tool.

Processes — the Linux term for programs running on the system — are shown in a table format. The default view shows all the processes that are currently being run by your user account. Alternatively, you can watch all the processes, or just the ones that are actively working, by clicking the View menu item from the menu bar and selecting which option you want to view. Selecting the Dependencies check box reorders the process list to show which processes started which other processes.

You can sort the table rows based on any of the table columns. For example, to see what processes are using the largest percent of CPU time, click the % CPU column heading. The System Monitor automatically sorts the rows based on percentage of CPU usage.

Other columns are available for you to add. Click Edit→Preferences to see the options available for the Processes tab. Table 17-2 lists the data available to display in the Processes tab table.

Table 17-2 System Monitor Process Table Columns	
<i>Column</i>	<i>Description</i>
Process Name	The program name of the running process.
User	The owner of the process.
Status	The status (either sleeping or running) of the process.
Virtual Memory	The amount of virtual system memory allocated for the process.
Resident Memory	The amount of physical memory allocated for the process.
Writable Memory	The amount of memory allocated to the process currently loaded into physical memory (active).
Shared Memory	The amount of memory shared between this process and other processes.
X-Server Memory	The amount of memory the process shares with the X Windows (GUI) server.
%CPU	The percentage of total CPU time the process is using.
CPU Time	The actual CPU time the process is using.

(continued)

Table 17-2 (continued)

<i>Column</i>	<i>Description</i>
Started	The time the process started running.
Nice	The system priority of the process. Higher nice numbers have lower priority on the system.
ID	The unique process ID (PID) the system assigned to the process.
Security Context	The security classification assigned by the SELinux security system.
Arguments	The name of the command and any command line arguments used to start it.
Memory	The amount of system memory the processes is using.

Here are a few other options you can set while you're in the Processes tab Properties dialog box:

- ✓ **Update Interval in Seconds:** How frequently the System Monitor refreshes the table data
- ✓ **Enable Smooth Refresh:** Gathers new process information before refreshing the table data, instead of refreshing table data as it gathers process information
- ✓ **Alert Before Ending or Killing Processes:** Provides a warning dialog box before you are allowed to terminate a running process “just in case”
- ✓ **Alert Before Hiding Processes:** Provides a warning when switching to hide a specific process

You can also control processes that you own directly from the Processes tab. Right-clicking a process produces a menu that allows you to stop, end, or kill a process, along with changing the priority of the process.

Resources

Clicking the Resources tab provides a quick overview of the hardware status of the system, as shown in Figure 17-15.




Figure 17-15:
The
Resources
tab in the
System
Monitor
tool.

There are three real-time graphs displayed:

- ✓ **CPU History:** Displays the running real-time CPU percent utilization. If the system contains more than one processor, each processor is shown as a separate history line.
- ✓ **Memory and Swap History:** Displays two running real-time graphs, one for the amount of memory used by the user, and one for the amount of swap space used by the user.
- ✓ **Network History:** Displays the amount of data send and received from the network interfaces.

The Resources tab can give you a quick overall picture of how your system is doing. If you see that the CPU or memory usage is running high, you can flip over to the Processes tab and sort the listing based on that parameter.

File Systems

Clicking the File Systems tab gives you a quick reference for the amount of disk space used on each mounted filesystem (shown in Figure 17-16) — see Chapter 7 for filesystems and mounting.

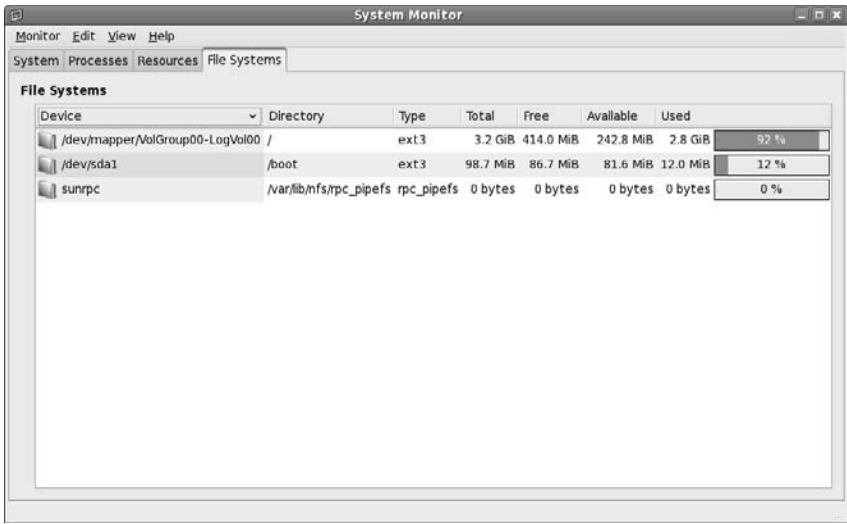


Figure 17-16:
The File
Systems
tab in the
System
Monitor
tool.

All hard drives and partitions are shown, along with their total disk space, free disk space, and used disk space. This allows you to get a quick indication of how much free space you have left on your hard drive.

As you can see, the System Monitor tool is really a bundle of various useful tools to help make system administration a little easier for you.

Chapter 18

A Secure Linux Box Is a Happy Linux Box

In This Chapter

- ▶ Implementing strong passwords
 - ▶ Keeping your system up-to-date
 - ▶ Plugging security holes
 - ▶ Using the System Logs Viewer
 - ▶ Securing your system by using best practices
-

I am Inspector Clouseau, and I am on official police business.

— Inspector Clouseau

You don't leave the front door of your house open when you go to work, do you? How about leaving it shut and locked but with a few nice, big windows open? The problem is that many people do this every day with their computers, and they don't even know it! In this chapter, I take a look at where your open doors and windows are and what you can do to secure them.

Every user's actions affect your overall system security. If your family members or officemates need access to your Linux machine, take the time to sit down and explain the facts of secure life to them. They can then apply this information to the other computers they use because these issues aren't specific to Linux.

Choosing Secure Passwords

The first line of defense from intruders is the passwords used on your system. For each account you have set up on your system, the passwords must be strong and difficult to figure out. If even one of the accounts has a weak password, you may be in for some trouble. Amazingly enough, in 70 percent of the cases where unauthorized individuals gained access to systems, the password for an account was the word *password* itself! When choosing good passwords, follow these rules:

- ✓ Don't use any part of your name.
- ✓ Don't use the names of friends, loved ones, or pets.
- ✓ Don't use birthdays, anniversaries, or other easily guessed dates.
- ✓ Don't use dictionary words.
- ✓ Don't keep your password written down near your computer, unless it's buried in something else, such as writing it into an address.
- ✓ Don't tell anyone your password. If someone needs to access specific files, give the person an account and set up permissions and groups properly so that they can do so.
- ✓ Do use a mix of lowercase letters, capital letters, and numbers.
- ✓ Do ensure that your password contains a minimum of eight characters.
- ✓ Do use acronyms made from sentences, such as having the password M8yodniT to stand for "My eight-year-old dog's name is Tabby."

Every person on your system needs to follow these rules, including you! Consider keeping a sheet of paper with these rules on it next to the machine.



I can't stress this advice enough: *Never* give out your password. Make sure that the people using your machine understand this rule. You can always find alternative methods to accomplish a task without giving out your password. If someone wants to use your machine, make an account for that person. Then they can have their own password!

Updating Software

All users can download and install new software. Of course, the programs they install are limited to the user's own permissions. The thing to be careful of here — as with any operating system — is that you don't get a version of a program that has been tampered with or is even an all-out fake trying to trick folks into installing it.

Most Linux applications and other Linux software programs are distributed by way of the Internet. In fact, the development cycle of new (and updates to) Linux software revolves around the Internet for file exchange, e-mail, and forum or newsgroup discussions. Make sure that you and other users of your Linux system are comfortable with the Web sites that are used and visited. You need to develop a *list* of trusted sites that provide you with the information you need and are not misleading in their presentation. As a starting point, you can *trust* all the Web sites referenced in this book because I have accessed them all. If either you or a user of your Linux system is unsure whether you can trust a particular Web site, do some research and perhaps ask others for their opinions.

Chapter 16 details how to keep your distribution and its software up to date. Please, please, please, do so! After all, as the person in charge, your job is to make sure that this computer stays intruder-free. In addition to making sure that you do all the same things a user would do for both your user accounts and the superuser (root) account, no matter which Linux distribution you're running, you must keep up-to-date with security problems.

Network holes


On a Linux server or workstation — or any computer at all, using any operating system — you should not have any network services running that you don't intend to use. Think of each network program running as a glass window or sliding glass door in your house. Each network service is a weak spot, and many nasty folks are out there on the Internet who like to go up to all the houses and make note of how many windows and glass doors are on them, what kinds they are, and how easy they are to breach.

Controlling your services

The more software you've installed on your system, the more services it may have running in the background. Fedora provides a utility for us to control what applications the system starts automatically. You can do this from the network service management program.

To open the network service management program, choose System⇨Administration⇨Server Settings⇨Services (see Figure 18-1). The Services option may be directly in the Administration menu: It depends on the collection of software you currently have installed. You'll also be asked for the root user password to ensure you have permissions to mess with the service settings.

Figure 18-1:
The Fedora
Service
Con-
figuration
dialog box.



The installed services are listed in the left side scroll window. Fedora starts the services with a check next to them at boot time. You can turn a service on or off simply by setting or removing the checkmark in a service's check box.

Services you may be interested in turning on or off include

- ✓ **apmd:** This service may not be necessary in anything but a laptop. It's used for monitoring battery power.
- ✓ **iptables:** This service is your firewall (more on the firewall in the section "Controlling and adjusting your firewall," later in this chapter). If you need to momentarily shut it down, you can do so using the Service Control dialog box.
- ✓ **ip6tables:** This service is for the new, improved IP version 6 protocol, which isn't being used in most of the world yet. You can most likely turn this off without breaking anything.
- ✓ **kudzu:** If you're using Fedora and keep getting bugged about hardware stuff at boot time, shutting off this service will stop those messages. You can run it manually as root if you change hardware later.
- ✓ **mdmonitor:** Shut this service off unless you implemented software RAID during your installation. (You had to go out of your way to do so, so if you don't know, you probably didn't!) If you change this service to on or off, make sure that mdmpd is also on or off (matching) as well.

- ✓ **pcscd:** The PC/SC smart card reader program. If you're not using a smart card to log into your Fedora computer, you can turn this off.
- ✓ **sendmail:** Even though you're probably not in need of a full-fledged mail server, shutting this service off can have unintended consequences because it's used to even handle internal mail on your system. Leave it on.
- ✓ **smartd:** If you're getting errors for this one at boot time, shut it off. It only works with certain IDE hard drives, so if you're not using that type of drive, it gives a (harmless) error.
- ✓ **spamassassin:** If you want to use this program in conjunction with your mail program, go for it! This program is used by default with Evolution in Fedora (see Chapter 10), so if you're using this combination of tools, leave this service on.
- ✓ **yum:** On Fedora, this program lets you run a nightly automatic update for those whose machines are connected overnight.



In Fedora, when you check or uncheck a service, you need to make sure that it does or doesn't turn on when you reboot. You need to use the Start and Stop buttons to start or stop it immediately. Use the bottom-right part of the dialog to see whether Fedora is running right now.

Controlling and adjusting your firewall

Even better (but just as essential) than turning off unnecessary services is to make sure that you have a firewall in place. A firewall is like putting a big bunker around your house. It would then have openings that only fit people wanting to do certain kinds of things. Friends could fit in through one door, family another, and package deliveries to another.

In computer networks, each of the services discussed earlier always comes in through the same door (*port*, in computer-world lingo). You use firewalls to prevent anyone from being able to so much as touch a door, or port, unless you've explicitly set it up so that they can do so. This technique is especially important if you're on a cable or DSL network (see Chapter 10), where there's always some overactive jerk out there using his computer to knock on every other computer on the network's doors to see where it can get in.

You probably already did some basic firewall setup during installation. If you ever want to make changes, choose System⇄Administration⇄Firewall and SELinux (you may be asked to enter your root password along the way). You should see the Security Level Configuration window, shown in Figure 18-2.

Figure 18-2:
The Fedora
Security
Level Con-
figuration
dialog box.



The Firewall Options tab has settings to easily configure your firewall options. You can enable or disable the entire firewall feature from the drop-down list box at the top of the window. If your computer is directly connected to the Internet — and most computers are — make sure to use Enable Firewall. The only time that you should not have this firewall in place is when your machine(s) are behind a strong firewall already, or you have a critical application that won't work otherwise. For just one application, though, that's one huge risk! You can find out how to open up the proper doors in the firewall for that one program instead.

Beneath the firewall setting drop-down list box is another list box containing the network services currently installed and available on your computer. Selecting a check box next to a network service enables remote connections to that service.

Select these items with extreme caution. Here's a list of the most common network services you'll see in the list:

- ✓ **HTTP and HTTPS:** Web stuff; only needed if you're running a Web server.
- ✓ **FTP:** FTP server stuff; you don't need it if you're not running an FTP server.
- ✓ **Mail (SMTP):** Allows remote mail servers to send mail directly to your computer. If you use an ISP for your e-mail (which most users do), you don't need to enable this.
- ✓ **NFS4:** Network File Sharing. Enable this if you want to share folders directly with other Linux or UNIX computers.

- ✓ **SSH:** Select this one to keep open. I explain it in the next section.
- ✓ **Samba:** Allows remote Microsoft computers to share files and printers on your Fedora computer. If you are connected to a local Microsoft network, you may need to enable this to share files. If you are directly connected to the Internet, it's not a good idea to have this enabled.
- ✓ **Telnet:** Please don't enable this one. It's used for remote interactive sessions with your computer, and is extremely unsafe in today's world.

The Secure Shell game (SSH)

One cool thing about Linux is that you can use the command line to connect to your account from anywhere, as long as you have the right software (and the machine you're connecting to isn't behind some kind of blocking software).



Some people tell you to use the `telnet` program to do this, but I beg you not to. Again, do not open the Telnet port in the security tool and do not use the `telnet` program. It sends information across the Internet in nice, raw text that anyone can snoop through.

First, you need to make sure that you have SSH installed, that you have an SSH server running, and that you enable SSH in your firewall. (See the section “Controlling and adjusting your firewall,” earlier in this chapter, to handle the third concern).

By default, Fedora installs both the SSH software and starts the SSH server to accept SSH connections from remote computers. If you don't want to provide this feature, you can stop the SSH service using the Service Configuration tool discussed earlier in the “Controlling your services” section.

If you elect to keep the SSH server running on your Fedora computer, you can use any SSH client software to connect to the Fedora computer from a remote computer. The software that you'll need depends on what type of computer you're trying to connect from.

Installing a Windows SSH program


If you want to connect to your SSH-enabled Linux box — or, actually, to any computer set up to accept SSH connection, not just a Linux one — from a Windows computer, you have several excellent choices of software packages available to choose from.

One popular SSH client package for the Windows world is PuTTY. This is a free software package released under the MIT license — similar to the GNU GPL license that most Linux distributions use — so you can use it to your heart's content without paying a dime.

The home page for PuTTY (<http://www.chiark.greenend.org.uk/~sgtatham/putty/>) provides a Download link to the download area. From here, you can select which version of PuTTY to use (you can even download the program source code and play with it yourself if that's what you're into). The easiest way to get PuTTY running on your Windows PC is to select the *putty.exe* Windows binary file for download. This is the complete PuTTY program in one simple file.

After downloading the *putty.exe* file, locate where you downloaded it on your computer by using your file manager program (such as Windows Explorer) and double-click it. The main PuTTY window, shown in Figure 18-3, appears.

Figure 18-3:
The PuTTY
main
window.



That's all there is to it: PuTTY doesn't require any installation setup. You're now ready to connect to a remote computer using SSH!


Setting up and making your SSH connection in Windows

Using PuTTY is almost as easy as downloading it. From the main PuTTY window:

1. Set the emulation features you want using the left side menu items.

As you select an item, the right side of the PuTTY window shows the possible configuration items related to it. For example, you can set the window appearance by clicking the Appearance entry under the Window section. Figure 18-4 shows the Windows appearance settings that appear.

Figure 18-4:
The PuTTY
window
appearance
configura-
tion.




2. Click Session in the left side menu to show the session options.
3. Type the hostname or IP address of the remote computer you want to connect to in the text box at the top of the window.
4. Ensure the SSH radio button under the address is selected (this should already be selected by default).
5. If you want to save the connection settings, put a name in the Save Sessions text box and click Save.
6. Click the Open button to start the SSH session.

You can select the saved session from the list box. PuTTY restores all of the configuration settings to what you had saved.

The terminal emulation window appears, asking you for the user id and password to use for the connection. After entering your login information, you should get a command line prompt for the remote computer, as shown in Figure 18-5.

Figure 18-5:
Connecting
to a remote
Fedora
computer
from a
Windows
PC.



Now you can enter Linux commands just as if you were on the Fedora console! SSH ensures that all of your session data is safely encrypted before it's sent out across the network. When you're finished, type **exit** at the command line and your session closes.

Connecting to your Linux box from another Linux box with SSH

Yes, you can connect from another Linux box, too. This task is a bit less complicated. Open a terminal window (see Chapter 6) and follow these steps:

1. Type `ssh username@ipaddress` or `ssh username@machine fullname` **to open the connection.**

For example, type `ssh dee@192.168.1.6` or `ssh dee@computer.example.com`. After you do this step, the following text appears:

```
The authenticity of host '192.168.1.6 (192.168.1.6)'  
can't be established.  
RSA key fingerprint is ed:68:0f:e3:78:56:c9:b3:d6:  
6e:25:86:77:52:a7:66.  
Are you sure you want to continue connecting (yes/no)?
```

2. Type **yes** and press **Enter**.

You now see these lines:

```
Warning: Permanently added '192.168.1.6' (RSA) to the  
list of known hosts.  
dee@192.168.1.6's password:
```

3. Enter your login password and press **Enter**. Now you're in!

Close the connection by logging out of the account (type **logout**).

Connecting to your Linux box from a Macintosh running OS X with SSH

The process from a Macintosh is similar to that under Linux. Go to Applications→Utilities→Terminal.app, which opens a command line window for you. Then type

```
ssh IPaddress
```

to access the same user account on the remote machine (again, you can use a computer's full name rather than the IP address), or type

```
ssh login@IPaddress
```

if you want to access the account *login* instead of the same account you're using on the Mac.

Software holes

When someone is already in your system — whether or not they're allowed to be there — you have additional security concerns to keep in mind. One of these involves what software you have on the machine. Believe it or not, each piece of software is a potential security hole. If someone can get a program to crash in just the right way, they can get greater access to your system than they should. That's a very bad thing!

One way to close software holes is to remove all programs you don't need. You can always add them later, if necessary. In Fedora, use the Add/Remove Software tool (see Chapter 16) to remove any unnecessary software from your computer.

If it turns out that, as a result of dependencies, you lose other software that you want to keep, make sure to cancel the removal.


Keeping an Eye on Your Log Files with the System Log Viewer

One other security issue you may want to configure concerns *log files*. Your network programs, kernel, and other programs all run log files, which contain records of what has been happening on your system. You may be amazed at just how much information gets put in them! They're mostly in `/var/log`; take a look sometime.

Fortunately, tools are available that can help mere mortals sift through the wheat to look for the chaff of bugs and intruders. Fedora provides the System Log Viewer, which allows us to easily scan through the major log files on the computer. Start the System Log Viewer by choosing System→Administration→System Log (see Figure 18-6). It'll ask you for the root password, as all of the log files are restricted for only root access.

The System Log Viewer provides quick links to the boot.log, maillog, messages, secure, and Xorg.0.log log files. If you want to view any other log files, you can do that by selecting File→Open from the menu bar, selecting the log file whose contents you want to view, and then clicking Open.

Figure 18-6:
The Fedora
System Log
Viewer.



As you can see, there are lots of log files to dig through. Here's a quick summary of what these log files contain:

- ✓ **messages:** This is the main system log file. It contains error and status messages from the operating system as well as many applications and services. All of the messages entered in the messages log file are in the format:

```
timestamp system application: message
```

You can use the filter feature to look at only messages from a specific application to help identify problems on the system. A sample entry looks like this:

```
Apr  9 19:34:55 testbox dhclient: DHCPREQUEST on eth0
to 192.168.73.254 port 67
```

This entry shows the eth0 network interface used the dhclient application to request an IP address via DHCP.

- ✓ **boot.log:** This log file contains messages generated from the boot-up process (most of which scroll by way too fast to read while you're booting the system). As applications and services start, they report their status in the boot.log file. The boot.log file uses the same format as the messages log file. This is a great place to look for answers if Fedora doesn't recognize any hardware on your system.
- ✓ **secure:** The secure log file reports every instance of a security request on the system, such as when a user tries to log in, or requests special privileges. Check this log if you can't log into the system with a user account.
- ✓ **maillog:** The maillog reports error and status messages from the mail system. Each time a message is sent or received, a log message is generated and stored in this log file. Check here if you can't send or receive mail messages.
- ✓ **X.org.0.log:** This log file tracks the status of the graphical X.org server on the computer. The X.org server provides the windows environment for the desktop. Check here if you're having problems with the desktop displaying on your monitor.

The secure log file contains information about connection attempts and failed access attempts. It's a good idea to peruse through this log file on a regular basis to watch for intruders. If you see a lot of denied access requests, you may have someone trying to hack into your system.

It also records each time a user requests root permissions. For example, when I started the System Log Viewer application, it required me to enter the root password. This generated an entry in the secure log file:

```
Mar 20 13:38:22 testbox userhelper: running /usr/sbin/  
gnome-system-log with root privileges on behalf  
of 'rich'
```

This log entry identifies the application that was run with root privileges, and the user id that requested it.

It's also a good idea to peruse through the messages log file on a regular basis. If any system services fail to start, you'll find out from entries in this log file.

Locating Security Resources

You can find a plethora of information on the Internet about desktop, network, and Linux security. Because of the massive volume of information available, I list some Web sites I like for security issues:

- ✓ **www.sans.org**: One of the major security-related sites on the Internet.
- ✓ **www.grc.com**: Provides some interesting tools, such as tools to test which ports are open on a system. Also, this site features many excellent articles dealing with system and network security. Click the Shields Up graphic.
- ✓ **www.tldp.org/HOWTO/Security-HOWTO/index.html**: The Linux Security *HOWTO*.
- ✓ **www.linuxsecurity.com/**: Presents a plethora of information from Linux Security.com.
- ✓ **www.securityspace.com/**: Has lots of information about security issues and tools for different operating systems.

Chapter 19

Ten Steps to Making Your Own Wireless Access Point

In This Chapter

- ▶ Hello, Master
 - ▶ Steps to making a wireless access point
 - ▶ Troubleshooting
-

You see, wire telegraph is a kind of a very, very long cat. You pull his tail in New York and his head is meowing in Los Angeles. Do you understand this? And radio operates exactly the same way: you send signals here, they receive them there. The only difference is that there is no cat.

— Albert Einstein, when asked to describe radio

The wireless phenomenon has hit just about everywhere. You can now change TV channels, answer the phone, or start your coffee maker from just about any place in the house. Wireless technology has also invaded the computer world.

Back in the old days (the early 2000s), you had to run cables all over your house in order to connect multiple computers to the Internet. Now you can use wireless technology to send and receive *packets* (chunks of data traveling over a computer network) from any computer without having to poke holes in any walls.

You can now purchase a *wireless access point* (WAP) from just about any electronics store. A WAP connects to your wired network, and then produces radio signals throughout your house (and even outside your house for a small distance). Computers with wireless cards can receive those signals and connect to the wired network.

As we mentioned in Chapter 9, Fedora supports many wireless cards as a client if you already have a wireless access point on your network. But what about if you have an Internet connection, but no wireless access point? Well, Fedora can come to your rescue there as well.

Hello, Master!

Every wireless network must have a WAP for remote computers to connect to. In wireless lingo, the WAP is called the *master device* — the central point where computers connect to interact with other devices on the wireless network. All network activity on the wireless network passes through the master, even if it's just between two computers. When you're configuring your Fedora computer's wireless card, you may come across the term *managed mode*. *Managed mode* means that your computer is connected to a master device on the network. This is demonstrated in Figure 19-1.




Figure 19-1:
Using a
master
wireless
device for
wireless
connectivity.

There's also *ad-hoc mode* — where two clients can talk directly between themselves without a master. However, this mode is limited to only two clients. For the purposes of this chapter, don't bother trying to use this method for your access point: The goal here is to support multiple wireless clients.

The key to creating your own wireless access point is to build a master device that all the computers in the house can connect to. When you purchase a commercial wireless access point, the manufacturer configures it to be a master, allowing remote wireless clients to connect to it. However, you don't need a commercial access point for your wireless network: You can accomplish the same thing with the wireless network card in your PC.

Although typical Linux setups don't usually use them as such, you can configure most computer wireless network cards to act as network masters. (Some don't support this feature, so you may have to check your wireless card's Web site to make sure.) The next sections walk you through the steps required to use your Fedora system as a wireless access point, using a wired connection to connect to the Internet.

Step 1: Plan Your Network

As with any other network implementation, it pays off to plan what you want to do before you start. The disadvantage to wireless networks is that they are extremely susceptible to environmental conditions. Wireless radio waves have limited range, and have difficulty penetrating some substances, such as walls made of metal, brick, or even solid wood.

Using a normal PC wireless network card as an access point aggravates these limitations. Most wireless network cards designed for PCs use a small antenna to be as unobtrusive as possible. Although this may look nice, it doesn't do much for wireless network range.

The first part of the plan is to try to find a network card that uses a large antenna, or has the capability of being connected to an external antenna. You don't need to set up an amateur radio antenna farm, but having something larger than a stub sticking out of the back of your PC is helpful.

Wireless access points also share a rule with real estate — location, location, location. Where you place your wireless access point is crucial to how well it will operate. Obviously, if you have a wired Internet connection, you need to locate your access point somewhere within reach of the Internet connection, but try to make that as centrally located as possible.

After you've decided on your location and antenna configuration, you're ready to start configuring your Fedora system.

Step 2: Configure Your Network Card

Before you can share your network connection wirelessly, you must be connected to the network. For this example, I'm going to assume you have a wired Internet connection (such as a cable or DSL modem). You must first connect your Fedora system to your network by using a wired network card (see Chapter 9).

Normally when you install Fedora it automatically detects your network card, and, if you're connected to a network, it also detects the network settings. Sometimes, however, you need to manually tweak things to make the card work on your network. Here's a summary of the steps from Chapter 9 that are required to do that:

1. From the GUI desktop, choose System⇨Administration⇨Network.

This starts the Network configuration tool. If you're not logged in as the root user, it asks you for the root user password to continue. When the Network configuration tool appears, it shows a listing of detected network cards. Note the device name of the network card (usually it's called `eth0`). You'll need that information later on in the process.

2. Double-click the network card entry.

The Ethernet Device dialog box (shown in Figure 19-2) appears, showing the current configuration of the network card.



Figure 19-2:
The
Ethernet
Device
dialog box in
the Network
configura-
tion tool.

The Ethernet Device dialog box contains several settings that you can alter:

- **Nickname:** Allows you to assign a different (possibly meaningful) name to the device rather than eth0 or eth1.
- **Activate Device when Computer Start:** Allows Fedora to automatically activate the network card at boot up without any intervention from you.
- **Allow All Users to Enable and Disable the Device:** By default, only the root user can start and stop the network card connection. Checking this check box allows any user to do that.
- **Enable IPv6 Configuration for This Interface:** This is a future standard for providing Internet addresses. At this time, most ISPs don't support IPv6, so leave this check box empty.

3. Select the Statically Set IP Addresses radio button to set a static IP address of 0.0.0.0.

When setting up a wireless access point, the Fedora system is used to pass network packets between the wired and wireless network cards. You won't need a network address to do that.

4. Click OK to save the settings.

The Network configuration tool saves the new network card configuration settings automatically. If the network card is not in Active status, click the Activate button to get it started.

You should now have your wired network card active, but you won't be able to use it on the network as you haven't assigned an IP address for it. If you're just using your Linux computer as an access point, you won't need to configure an IP address. Next up is the wireless card.

Step 3: Install Your Wireless Card

Installing and configuring a wireless network card in Linux can be a tricky thing. This is one area where many Linux distributions fail miserably.

Linux supports only a limited number of wireless network cards. Actually, it's not so much the cards that aren't supported as the chips used inside the cards. One of the biggest headaches in the wireless network card world is that manufacturers often switch the types of chips used in the same model card in different manufacturing runs (revisions). A Linux driver that works fine for card A, revision 1 may not work at all for card A, revision 2.

If Fedora detected your wireless network card during installation, you can use the standard Network Configuration tool to configure it (we walk you through that process in a little bit).

If your wireless network card wasn't detected, or if you try to install a new wireless network card, you have a few options available to you:

- ✓ Use the Network Configuration tool to select the new card
- ✓ Use the Linux tools provided by the card manufacturer
- ✓ Use the ndiswrapper Linux application

First, try using the Network Configuration tool to see if Fedora recognizes your wireless card and has a driver for it. This process is detailed in Chapter 9. Just follow those steps, and then jump back here if you get your wireless network card working.

If you can't find your network card in the Network Configuration tool list, you'll have to do some experimenting. The next best thing is if your wireless network card manufacturer provides a Linux tool to help install the card. Check your wireless card's supplied drivers CD and its Web site. Hopefully, you'll be able to find a Linux driver or tool for your particular card. If you do, follow the instructions provided by the manufacturer.

If you still can't get your wireless card to work, there's one more chance for you. The ndiswrapper application is an open source project that attempts to create a Linux driver from the Windows drivers supplied with the wireless network card.

The ndiswrapper Wiki site (http://ndiswrapper.sourceforge.net/mediawiki/index.php/Main_Page) provides information on which wireless network cards work with ndiswrapper. Check the Wiki site first to see if there's any hope of using your particular wireless network card. Lots of people post solutions to their wireless network card problems, so there's a good chance you'll find information there about your card. If your wireless network card doesn't appear in the list, you are officially out of luck and need to buy a new card.

If you find that you can use ndiswrapper to support your wireless network card, you can continue with this section. You must first install the kmod-ndiswrapper package:

1. Click Applications⇨Add/Remove Software.

This starts the Package Manager tool. If you're not logged in as the root user, you'll be asked for the root password to continue.

2. **Click the Search tab to search for a software package.**
3. **Type `kmod-ndiswrapper` in the text box and click the Search button.**
When the search completes, you'll see the current `kmod-ndiswrapper` package listed in the results window.
4. **Select the `kmod-ndiswrapper` entry, and then click Apply to start the installation.**

After you finish with the software installation, you need to load the `ndiswrapper` driver:

1. **Copy your wireless network card's Windows drivers to a folder in your home directory.**
Often the drivers available on the CD that comes with your network card are old. Try visiting your wireless network card's Web site to see whether new drivers are available.
2. **Click Applications→System Tools→Terminal.**
The Terminal tool appears, showing a command line prompt.
3. **If you are not already logged in as the root user, you must become the root user.**

To do that at the command line, type `su` and then enter the root password. The command line prompt changes to a pound sign, indicating you are now the root user.

4. **Change to the folder where you copied the Windows drivers.**

Use the `cd` command line command to change to the folder name where you copied your drivers:

```
cd drivers
```

5. **Install the `.inf` file from the Windows drivers.**

Use the `-i ndiswrapper` option to install the `.inf` file provided by the wireless card manufacturer:

```
ndiswrapper -i file.inf
```

6. **Install the `ndiswrapper` module.**

Use the `-m ndiswrapper` option to install the `ndiswrapper` module:

```
ndiswrapper -m
```

7. **Reboot the system.**

After the system reboots, follow the normal steps to install a wireless network card as shown in Chapter 9. The `ndiswrapper` driver appears as a new wireless network card called `wlan0`.

If all is well and Fedora recognizes that your wireless network card exists, it's time for you to configure it to be a wireless access point. Here's what you need to do to change that.

1. From the GUI desktop, choose System⇨Administration⇨Network.

If you're logged in as a regular user (not root), you're prompted to enter the root password. The Network Configuration dialog box opens. Note the name of the wireless network card (if it's the second installed card, it'll be called `eth1`). If you used the `ndiswrapper` technique to install your wireless network card, it'll use an alias name called `wlan0`.

2. In the Network Configuration dialog box, double-click your detected wireless network card.

The Wireless Device Configuration dialog box appears (as shown in Figure 19-3).



Figure 19-3:
The
Wireless
Device Con-
figuration
dialog box in
the Network
Configura-
tion tool.

3. Click the Wireless Settings tab, and then enter your wireless access point configuration settings.

You must configure several wireless settings to connect to your access point:

- **Mode:** Determines if the wireless card is managed by an access point, or is the master of other wireless cards. Use the drop-down list box to set the card to Master.
- **Network Name (SSID):** Wireless access points are identified by a Service Set Identifier (SSID). Enter a unique SSID for your network in this field. This can be any combination of letters, numbers, and symbols, up to 32 characters.
- **Channel:** You must select a specific wireless channel to use. There are 14 channels to choose from. Most wireless clients scan all channels looking for your SSID, whereas others must be specifically configured for the correct channel.
- **Transmit rate:** Some wireless network cards allow you to specify the transmit rate, whereas others autodetect the rate used by the remote access point.
- **Key:** Encrypting wireless access points is a necessity in today's networking environment. This field allows you to enter the encryption key used by your access point. However, to get started, leave this field blank for now.

For the time being, create the wireless access point without an encryption key. After you know the wireless access point is working, you can play around with adding an encryption key. Starting off with one can often make troubleshooting more difficult than it needs to be.

4. Click the General tab to remove the IP address for the wireless card.

Set a static IP address of 0.0.0.0 for the wireless network card. Because you're using the card as a bridge, you don't need an IP address for the card.

5. Click OK to finish configuring the wireless network card.

After the wireless network card appears in the Network Configuration list, you can easily start and stop it by using the Activate and Deactivate buttons.

Step 4: Install Bridging

For your next trick, you need to build a bridge between the wired network card and the wireless network card. No, not a suspension bridge, but a network bridge.

Network bridges are used to bridge a gap in the network. They connect two network systems together on the same router subnet. Routers are used to move traffic between two dissimilar networks, whereas bridges move traffic between to similar networks that must be joined together.

The bridging software required for this project is already prepackaged in a Fedora software package. The `bridge-utils` package contains all the software components you need to build your network bridge. All you need to do is install the package:

- 1. Click Applications → Add/Remove Software.**

This starts the Package Manager tool. If you're not logged in as the root user, you'll be asked for the root password to continue.

- 2. Click the Search tab to search for a software package.**

- 3. Type `bridge-utils` in the text box, and click the Search button.**

When the search completes, you'll see the current `bridge-utils` package listed in the results window.

- 4. Select the `bridge-utils` entry, and then click Apply to start the installation.**

You've now completed the software installation: next, you need to use the bridging software to create the bridge.

Step 5: Build a Bridge

Building the bridge is possibly the most complicated of all the steps. Unfortunately, there aren't any fancy graphical tools to help you out here. You have to dig into the Fedora command line and do things the old-fashioned way — typing commands one by one at the command prompt.

The Linux Ethernet bridging command you need to use is `brctl`. This command allows you to create a network bridge, add network interfaces to the bridge, and monitor the bridge.

The format of the `brctl` command is

```
brctl [command]
```

where *command* is the bridging command to execute. The `brctl` commands you'll use are shown in Table 19-1.

Table 19-1 The brctl Commands

<i>Command</i>	<i>Description</i>
addbr	Create a new network bridge
delbr	Remove an existing network bridge
addif <i>br if</i>	Add network interface <i>if</i> to bridge <i>br</i>
delif <i>br if</i>	Remove network interface <i>if</i> from bridge <i>br</i>
stp <i>br</i> state	Set the spanning tree protocol for bridge <i>br</i> to on or off

To execute command line commands, you'll need to open a Terminal session:

1. Click Applications⇨System Tools⇨Terminal.

The Terminal tool appears, showing a command line prompt.

2. If you are not already logged in as the root user, you must become the root user.

To do that at the command line, type **su**, and then enter the root password. The command line prompt changes to a pound sign (#), indicating you are now the root user.

3. Create a new bridge instance called br0.

Use the **brctl** command to create a new bridge:

```
brctl addbr br0
```

This command creates a new bridge called br0.

4. Temporarily disable the Ethernet interfaces.

Before adding the interfaces to the bridge, you need to stop them from receiving network traffic. Do this by using the **ifconfig** command:

```
ifconfig eth0 down  
ifconfig eth1 down
```

5. Add the network cards to the newly created bridge.

Again, use the **brctl** command to add the wired network card by using the interface name found in Step 2, and the wireless network card by using the interface name found in Step 3:

```
brctl addif br0 eth0  
brctl addif br0 eth1
```


That completes the bridge-building process. You can look at your new bridge by using the `show` command option:

```
# brctl show
bridge name      bridge id        STP enabled      interfaces
br0              8000.00402b2a3175  no              eth0
                eth1
```

The `brctl` command shows that the bridge is OK, and that it contains both of the network interfaces. Now you can activate both of the network interfaces and the new bridge by using the `ifconfig` command:

```
ifconfig eth0 0.0.0.0 up
ifconfig eth1 0.0.0.0 up
ifconfig br0 up
```

There should now be network traffic crossing between the two network cards.

Step 6: Save Your Hard Work

After typing all those commands at the command prompt, you've got a working bridge between the wired network card and the wireless network card. The only problem is all that work goes away when you reboot the system. Ouch.

Instead of having to type all the bridging commands every time you boot your system, the smart thing to do is to get the commands to run automatically when the system boots. The easiest way to do that is to place the commands in a script that runs at boot time.

The most likely candidate for that is the `/etc/rc.local` file. This file contains script entries that are run at the end of each boot up.

Here's how to get your commands into that script:

- 1. Log into the Fedora system as the root user.**

You need root permissions to modify the `rc.local` file, so the easiest way to do that is to just log in as root.

- 2. Click Applications⇨Accessories⇨Text Editor.**

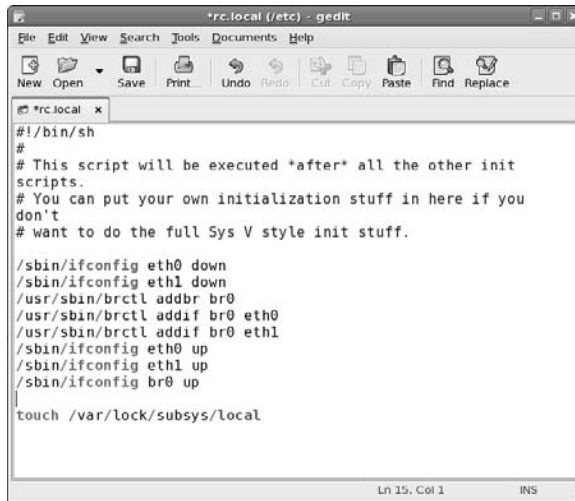
This command starts the `gedit` text editor program.

- 3. Click the Open icon in the toolbar, find the `/etc` folder, and then select the `rc.local` file to open.**

4. Add the `brctl` command line entries typed in Step 5 of the preceding step list to the file.

When referencing the `brctl` and `ifconfig` commands in the `rc.local` file, use the full pathname to the commands (`/usr/sbin/brctl` and `/sbin/ifconfig`). The resulting file should look like the one shown in Figure 19-4.

Figure 19-4:
Adding the
`brctl`
commands
to the `/etc/
rc.local` file.



```
*rc.local (/etc) - gedit
File Edit View Search Tools Documents Help
New Open Save Print Undo Redo Cut Copy Paste Find Replace

#!/bin/sh
#
# This script will be executed *after* all the other init
# scripts.
# You can put your own initialization stuff in here if you
# don't want to do the full Sys V style init stuff.

/sbin/ifconfig eth0 down
/sbin/ifconfig eth1 down
/usr/sbin/brctl addbr br0
/usr/sbin/brctl addif br0 eth0
/usr/sbin/brctl addif br0 eth1
/sbin/ifconfig eth0 up
/sbin/ifconfig eth1 up
/sbin/ifconfig br0 up

touch /var/lock/subsys/local

Ln 15, Col 1  INS
```

5. Click the Save icon in the toolbar to save your changes.
6. Exit the gedit tool.

That completes the process. Reboot your system, and everything should come up as expected. You can use the `brctl` with the `show` command to see if your bridge started correctly.

Step 7: Test Connectivity

The new wireless access point should now be on the air. To test this out, go to another wireless client and browse the network for new access points.

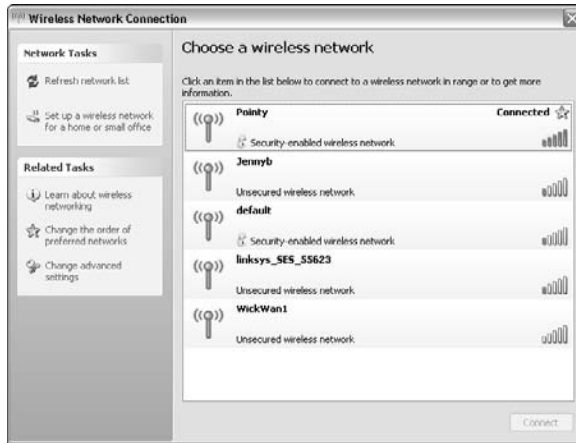
On a Microsoft Windows XP Home workstation, you can do that by using these steps:

1. Click Start → My Network Places.
2. Select View Network Connections from the left panel.

3. Find your wireless network card in the device list, and then right-click the icon. Select **View Available Wireless Networks** from the menu.

A list of the wireless networks your client sees appears, as shown in Figure 19-5.

Figure 19-5:
Wireless
networks
list for a
Microsoft
Windows
XP Home
workstation.



The new wireless access point (called Pointy in this example) should appear on your list of available networks.

4. Select the new wireless access point by double-clicking the entry.

The wireless network card should connect to the access point. If your network uses DHCP to assign network addresses, the client receives an address and should be good to go. If your network uses static IP address, you need to assign an appropriate address to the wireless network card on your client.

Congratulations: You now have a working wireless access point for your network. The next step is to go back to the wireless network card configuration and fix your little security problem.

Step 8: Configure Security

Okay, so now you have your wireless access point working, but there's just one small problem — anyone can use it! If you're running a public access network, I suppose that could be a good thing, but for most of us, this means problems.

The smart thing to do is lock down the wireless access point so only authorized devices can use it. Here's how:

1. From the GUI desktop, choose System⇨Administration⇨Network.

If you're logged in as a regular user (not root), you're prompted to enter the root password.

2. In the Network Configuration dialog box, double-click your detected wireless network card.

The Wireless Device Configuration dialog box (refer to Figure 19-3) appears, showing the current settings for the wireless network card.

3. Click the Wireless Settings tab to change your wireless access point configuration settings.

4. Enter a 13-character hexadecimal or ASCII encryption key in the Key text box.

The wireless network card uses the Wired Equivalent Protocol (WEP) to encrypt data sent across the wireless network. WEP requires a 13-character hexadecimal or ASCII key. Select a unique key to encrypt your data. If you select a hexadecimal key, you must precede the key with a 0x.

5. Click OK to exit the Wireless Device Configuration dialog box, and then click File⇨Save to save the changes.

Your wireless access point should now be secure. Not only must remote clients know the key to connect to the access point, but it will also be more difficult (but not impossible) for any snoopers to decode your wireless network traffic.

The next step is to get the remote clients working.

Step 9: Configure Remote Devices

Now that you've got your wireless access point encrypted, you must reconfigure your remote clients. Here are the steps to do that on a Microsoft Windows XP Home PC:

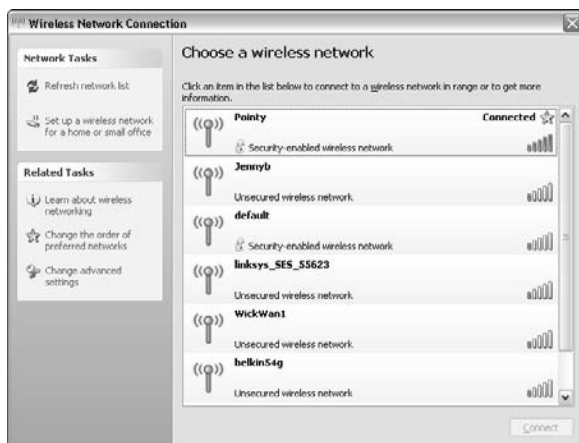
1. Click Start⇨My Network Places.

2. Select View Network Connections from the left panel.

3. Find your wireless network card in the device list, and then right-click the icon. Select View Available Wireless Networks from the menu that appears.

A list of the wireless networks your client sees appears, as shown in Figure 19-6.

Figure 19-6:
Looking
for an
encrypted
wireless
network in
Microsoft
Windows
XP Home
Edition.




The new wireless access point (called Pointy in this example) should now appear in your list of available networks with a small lock symbol. This indicates that the access point is using encryption.

4. Select the access point from the list.

This time, when you select the access point, a dialog box appears asking you to enter the encryption key, as shown in Figure 19-7.

5. Click OK to accept the encryption key and connect to the wireless access point.

Figure 19-7:
Entering the
encryption
key to
connect
with an
encrypted
wireless
network in
Microsoft
Windows
XP Home
Edition.



The wireless network card should connect to the access point. Again, if your network uses DHCP to assign network addresses, the client receives an address and should be good to go. If your network uses static IP addresses, you need to assign an appropriate address to the wireless network card on your client.

Step 10: Test Remote Devices

After you are connected to the wireless access point, the bridge created on the Fedora system passes all network packets from your wireless clients to the wired network.

Test this out by trying to access an Internet location with a Web browser. If you can get out onto the Internet, congratulations! You've built your own secure wireless access point.

In Case of Difficulty . . .

As you can imagine, with a system this complex, there's plenty of places for things to go wrong. The main thing to remember when troubleshooting problems is to break the system down into the individual components and test each one.

Here are a few things that could go wrong, along with tips on what to check for:

- ✓ **You don't see the wireless access point in the client's list of access points:** Check to make sure that the client is in range of the wireless access point. Remember, if you're using a PC wireless network card, the range may be somewhat limited.

Next, use the Network configuration tool to ensure that the wireless network card is active. If something goes wrong with your `rc.local` script, the wireless cards may not be activated.

- ✓ **You can connect to the wireless access point, but you can't connect to the Internet:** Check that the wired network card is active. Use the Network configuration tool to see the status of the wired network card, and activate it if necessary.

Also, use the command line `brctl` command to check the status of the bridge. You can use the `showmacs` command to see if the client's network card Ethernet address appears on the bridge's list of addresses.

- ✔ **Everything worked great when you manually set it up, but after rebooting, nothing works:** Check your `/etc/rc.local` script to ensure that you entered the `ifconfig` and `brctl` commands in the correct order. It's also possible that these commands were run before the wireless network card was ready and they failed.

Some PCMCIA and USB wireless network cards must be activated by the Fedora system before you can configure them. Sometimes the cards are still starting up when the `rc.local` script is run. If this is the case, you may have to use the Linux `sleep` command to place a short delay before you run the `ifconfig` and `brctl` commands.

Chapter 20

Ten Steps to Setting Up a Samba Print Server

In This Chapter

- ▶ Dancing the Samba
 - ▶ Setting up a print server
 - ▶ Troubleshooting
-

Making duplicate copies and computer printouts of things no one wanted even one of in the first place is giving America a new sense of purpose.

— Andy Rooney

Networks have revolutionized the computer world. Not only do they help you view fancy Web sites, they also give you the ability to share a single printer between two computers on opposite sides of the house. Being able to share a single printer among all of your household computers can save you lots of money.

Linux computers can share printers among themselves just fine. Chapter 17 shows how to map a Fedora computer to just about any type of network printer available. However, most likely you also have one or more Microsoft Windows systems on your network. Getting these to work with Linux printers is not always as easy.

This chapter walks through the steps necessary to make the printer connected to your Linux computer available to your Microsoft Windows computers on the network.

Dancing the Samba

Making your Linux computer talk with Microsoft Windows workstations and servers requires a special software package. Microsoft uses a proprietary (secret) network protocol to communicate between Windows devices on a network.

Fortunately, some enterprising open source enthusiasts worked to reverse engineer the Microsoft networking protocols and produced software to duplicate them. The result of this effort is the open source package called Samba.

Fedora includes the Samba software as part of the core Fedora software packages. I use the Samba software package to make any printers attached to your Linux computer visible to the Microsoft Windows workstation or server on the network.

The following steps guide you through the processes required to install and configure Samba and Linux to create a print server that everyone can use.

Step 1: Plan Your Sharing

Before you can start loading and configuring software, it's a good idea to get a handle on your network environment. Determining exactly what you're doing before you start doing it is always a good thing.

First, let's make sure we're talking about the same thing here. The network-printing world has gotten confusing. There are special types of printers available on the market called *network printers*. You can connect these types of printers directly to your network without using an intermediate computer.

Network printers are often heavy-duty printers, capable of printing thousands of pages a day. These printers are designed for medium- to large-office use and are usually too expensive for the casual home or small business user. These are not the types of printers we're talking about here.

This chapter is all about sharing a common, ordinary printer that you plug into your Fedora Linux computer or server. These are called *local printers* because you connect them directly (or locally) to a specific PC.

The beauty of sharing local printers is that you don't need to buy a separate printer for each computer on your network. Knowing how to convert a simple inkjet printer into a shared network printer is a great bargain.

If you haven't purchased a local printer yet, check which ones are compatible with Linux. Unfortunately, most printers don't come with Linux drivers. A great resource is the OpenPrinting Web site (www.linux-foundation.org/en/OpenPrinting). This site contains a database of all printers that are known to have Linux drivers, and even has a download section where you can download drivers for your printer if they aren't already included in Fedora.

Step 2: Connect a Local Printer

After you've got your printer set, you'll want to connect it to your Linux system. The two most common methods used to connect local printers are

- ✓ **Parallel port:** Using a special 25-pin printer cable
- ✓ **USB port:** Using a standard USB cable connector

Most modern printers now use USB cables. The USB system provides many benefits that were missing from the older parallel printer cable method.

When you plug a printer into your Fedora Linux system with a USB cable, Fedora automatically detects the printer make and model and creates a printer device for you. You can start printing without doing any work of your own.

However, if you plug a printer into your Fedora Linux system with a parallel printer cable, nothing happens. You have to manually configure the printer in the system before you can use it (ugh).

In either case, it's a good idea to look at the system printer configuration first to ensure that your printer is set up properly. That's the next step.

Step 3: Manually Add a Local Printer

Fedora uses the Printer Configuration tool to add and manage all printers. If you are using a parallel printer cable, you have to manually add your printer in the Printer Configuration window. Here are the steps for doing this:

1. Select System⇨Administration⇨Printing from the panel menu.

This starts the Printer Configuration window, as shown in Figure 20-1.

This requires root permissions to run. If you're not logged in as the root user, it asks you for the root password before starting.

Figure 20-1:
The Fedora
Printer Con-
figuration
window.



2. Click the New Printer button.

The New Printer Wizard starts. You must assign the new printer a name that the server will use both on the local system and to advertise the printer on the network. You also have the option of defining a description and the location of the printer. In a network environment, the server advertises these items along with the printer name. This can help users find the right printer if you have more than one.

After you have finished typing in your entries, click the Forward button to move on to the next wizard window.

3. Select the connection method.

For a parallel printer, select the LPT#1 option and click the Forward button.

4. Select the printer driver.

This can be the most frustrating part of the process. If your printer manufacturer is listed in the provided list, select it, and continue on by clicking the Forward button.

If your printer manufacturer is not selected, you may be out of luck. Sometimes you can find the appropriate Postscript Printer Description (PPD) driver on a manufacturer's Web site, or the OpenPrinting Web site (www.linux-foundation.org/en/OpenPrinting). If so, select that option in the wizard window and use the text box to find the PPD file on your system.

5. Select the printer model.

If you selected a printer manufacturer, the next step is to select the specific printer model from the list. If the model of your printer isn't listed, you'll have to try to download the appropriate PPD file and use that method.

Some models have more than one PPD driver file available. In this case, usually one is marked as being recommended. Try that first. When you've selected the appropriate printer model and driver, click the Forward button.

6. Finish the install.

The final window summarizes your selections. If you need to change anything, you can click the Back button to return to the appropriate wizard window. If everything's OK, click Apply to create the printer definition.

After you've installed the printer, or if Fedora automatically detected your USB printer, it's a good idea to check over the printer configuration items. That's brings us to the next step.

Step 4: Set the Printer Configuration

The Printer Configuration window is divided into two panes. The left pane lists the items that can be configured. The right pane is used for showing configuration settings.

By default, the left pane of the window shows two general items:

- ✓ **Server Settings:** Allows you to set options that apply to the print server
- ✓ **Local Printers:** Allows you to set options that apply to the individual configured printers

As you click on an item in the left pane, the tool presents information in the right window. Here are the configuration items necessary for the print server:

1. Click the Server Settings item.

When you click Server Settings, it displays a set of check boxes for items to configure:

- **Show Printers Shared by Other Systems:** Scans the network for advertised printers and adds them to the list of available printers
- **Share Published Printers Connected to This System:** Advertises all local printers marked as shared on the network

- **Allow Remote Administration:** Allows remote connection to the printer configuration system
- **Allow Users to Cancel Any Job (Not Just Their Own):** Lets others cancel print jobs sent to the printer but waiting in the print queue
- **Save Debugging Information for Troubleshooting:** Instructs the printing system to create debugging files for observing what's happening

2. Make sure that the Share Published Printers Connected to This System check box is checked.

By checking this setting, Samba automatically shares any printer created on the system. (Don't worry if there are any individual printers you prefer not to share, you can configure them to not be shared.)

3. Click the specific printer entry under the Local Printers section in the left window pane.

The Local Printers option provides a list of all automatically and manually installed printers. Clicking the entry produces the Printer Configuration Details window, as shown in Figure 20-2.




Figure 20-2:
The Printer Configuration Details window.

4. Check the Shared check box in the right window pane.

You can also set various other options (described in detail in Chapter 17) for the printer, such as setting a banner page, restricting the printer to specific users, and setting the page size and paper source for the printer.

5. Click the Print Test Page button.

It's always a good idea to make sure the local printer is working properly before you try printing from a remote client.

Step 5: Install Samba

To share your printer with Microsoft Windows workstations and servers, you must install the Samba software package. Samba allows your Linux system to communicate with Windows devices by using the Microsoft networking protocol.

As shown in Chapter 16, you can easily install Samba by using the Fedora Package Management system. Here's how:

1. **Click Applications⇨Add/Remove Software.**

This starts the Fedora Package Management system. You can easily add or remove software packages from here.

2. **In the Package Management window, select the Servers option on the left side, and the Windows File System option on the right side.**

The Windows File System option includes one optional package — `system-config-samba`. Ensure that you select this package for installation as well. It will make configuring the Samba software easier.

3. **Click the Apply button on the lower right to start the software installation.**

Step 6: Configure Samba

After you install Samba, you must configure it for your specific Microsoft network environment. Microsoft networking uses two different methods for joining network devices together into groups:

- ✓ **Workgroup:** Defines a small number of workstations and servers that belong to an ad hoc network group
- ✓ **Domain:** Defines a larger number of workstations and servers that belong to a controlled networking environment

Because we are working with a small network environment, I assume that you're using a Microsoft workgroup on your network. In a workgroup environment, you must configure each workstation and server to belong to a specific workgroup name. For the network devices to see your shared printer, you must add your Linux server to the workgroup:

1. **Select System⇨Administration⇨Server Settings⇨Samba from the panel menu.**

This starts the Samba Server Configuration window (if you're not logged in as the root user, you'll be asked to supply the password for root first).

2. Click Preferences⇨Server settings from the menu bar.

The Server Settings configuration window, shown in Figure 20-3, appears.

Figure 20-3:
The Samba
Server
Settings
con-
figura-
tion
window.



3. In the Server Settings window, enter the workgroup name for your Windows network in the Workgroup text box.

By default, Fedora uses the Linux hostname of your system as the Windows network device name on the network. You can also provide a brief description of your host in the Description dialog box provided. Network devices see the description along with your hostname when they browse the network.

Click OK when you are done, and then close the Samba Settings window.



If you enabled the SELinux security software during installation, you'll need to tell it to allow Samba to operate. To do that, select System⇨Administration⇨SELinux Management to start the SELinux Management tool. From there, select Boolean, and then click the Samba drop-down list box. Select the Allow Samba to Share Any File/Directory Read/Write check box. That's all there is to it!

Step 7: Start Samba

Unfortunately, Samba does not start by default when you install it. You can elect to either manually start Samba when you want to, or have the Linux system start it automatically when it boots.

For either option, you must use the Fedora Service Configuration window:

1. Select System⇨Administration⇨Server Settings⇨Services from the menu.

This starts the Service Configuration window. Again, you must be the root user to use this tool, so if you're not already logged in as root, you'll be asked for the root password to start. When the Server Configuration window starts, you see all the services that are available on the Fedora system, as shown in Figure 20-4.

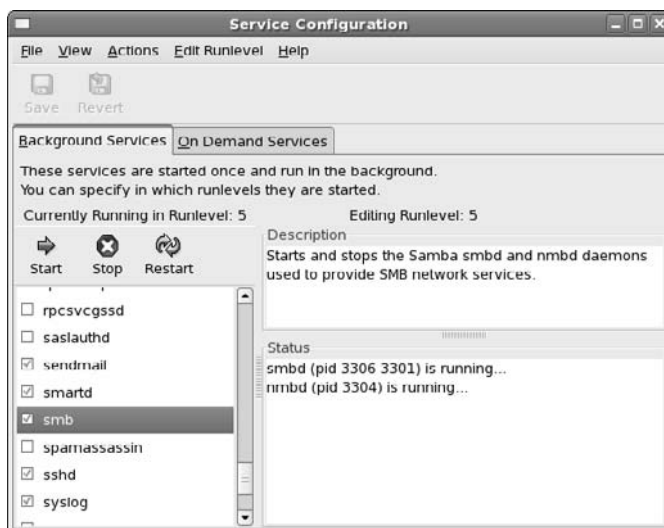


Figure 20-4:
The Fedora
Service
Con-
fig-
uration
window.

The server automatically starts services with a green check mark in the check box when the system boots. You can also start, stop, or restart individual services by using the appropriate button.

The Services Configuration window labels the Samba service as **smb**. This service actually controls two separate Samba services:

- **smbd**: Controls client connections to the Samba server
- **nmbd**: Controls advertising shared devices on the network

2. Select the **smb service and click **Start**.**

When you start the **smb** service, both of the Samba servers start. You see the status of the services in the right window pane.

3. If you want Samba to automatically start at boot up, select the check box next to **smb.**

Step 8: Open the Firewall Port

Now that the Samba server is running, you need to make sure remote network clients can connect to it. Because, by default, Fedora is a secure system, all unused network ports are blocked by the firewall.

The final step is to open the Samba server port so remote client packets will be accepted by the firewall and passed through to Samba:

1. Select **System**⇨**Administration**⇨**Firewall & SELinux** from the menu.

This starts the Security Level Configuration window, as shown in Figure 20-5.




Figure 20-5:
The Fedora
Security
Level Con-
figuration
window.

You must be the root user to have access to this tool, so if you're not logged in, you must enter the root password. This tool has three tabs:

- **Firewall Options:** For opening specific application ports
- **Advanced Options:** For opening specific TCP ports
- **SELinux:** For controlling how Linux locks specific applications

2. In the Firewall Options area, select the Samba entry.

The Firewall Options show a list of common services that are configured on the system. Selecting the Samba entry marks it as a trusted service and allows network clients to connect to it.

Step 9: Let Guests Print

By default, Samba uses the standard Windows network logon accounts to control network security. However, for a print server on a small network, this can be overkill.

For small networks, you most likely want to let anyone on the network print. In such a situation, it's kind of silly to create unique user accounts for everyone. Instead, you should just allow anyone to print on the shared printer.

Unfortunately, the Samba Configuration window doesn't provide for setting this feature, so you'll have to dive into the Samba configuration file and change it yourself. Here's how to do that:

- 1. Log in to the Fedora system as the root user.**

The Samba configuration file is protected so only the root user can edit it.

- 2. Click Applications→Accessories→Text Editor.**

- 3. Click the Open icon and browse to the Samba configuration file, /etc/samba/smb.conf.**

- 4. In the configuration file, look for the [printers] section.**

The [printers] section contains all the Samba settings for controlling printer sharing. You should see the following entries:

```
comment = All Printers
path = /usr/spool/samba
browseable = no
# Set public = yes to allow user 'guest account' to
    print
; guest ok = no
; writeable = no
printable = yes
```

- 5. Remove the semicolon in front of guest ok and change the no to a yes.**

This enables guests to access and use the shared printers.

- 6. Save the configuration file in the same location and exit the text editor.**

Step 10: Map a Network Printer

The last step is to go to a Microsoft Windows workstation or server and map to the shared printer on your Linux system. If you have a Windows XP Home Edition workstation, you can do this by following these steps:

1. **Select Start→Control Panel.**
2. **Select Printers and Other Hardware.**
3. **Select Add a Printer.**

This starts the Add Printer Wizard. Click Next to continue to the next window.

4. **Select the A Network Printer, or a Printer Attached to Another Computer radio button and click Next.**
5. **Select the Browse for a Printer radio button.**

If you know the printer and print server name, you can also manually enter it by using the format `\\servername\printername`.
6. **The Shared Printers text box shows a list of detected print servers on the network. Double-click your Linux server name and select the shared printer you want to map to.**

The Linux print server uses the same name as the hostname of the Linux system. After selecting the desired shared printer, click Next. A couple of Windows warning boxes appear warning you about the dangers of using network printers, but hey, this is your own network printer, so don't worry about the warnings.

7. **Select the appropriate printer driver from the manufacturer and model lists. Click OK when finished.**

Troubleshooting

In an ideal world, things should be working just fine for you. You should be able to map any Windows workstation or server on your network to the new Linux print server, and print to any locally connected printer on the server.

However, things don't always go as planned. There are plenty of opportunities for things to go wrong in the network world. Here are a few things to watch out for if things aren't working with your print server:

- ✓ **Test the local printer:** We can't stress this step enough. We've run into plenty of situations where Administrators have spent countless hours trying to troubleshoot a network printer problem only to eventually find out that the local system can't even print.

Make sure that you can send print jobs from the Linux system to the local printer, and that the print job comes out OK. Often a mismatched PPD driver can cause odd behavior with printers.

- ✓ **Remember to allow guest printing:** By default, Fedora configures Samba to block guest printing and doesn't provide an easy way to enable it. If you don't manually edit the `smb.conf` file to enable guest printing, no one is able to use the shared printer (which kind of defeats the purpose of sharing it).
- ✓ **Check both global and local printer sharing:** The Fedora Printer Configuration window provides two locations that must be set for each printer. The global Server Settings section must be set to enable sharing published printers, and the Local Printer must be set for sharing. If either of these settings is not set, you won't see the printer when browsing the network.
- ✓ **Check network connectivity:** Finally, the bane of many network problems is, of course, the network. Check to make sure that you have proper network connectivity from your Linux system. Try pinging the IP address of the Linux system from a Windows workstation or server on the network.

Chapter 21

Ten Steps to Setting Up a File Server

In This Chapter

- ▶ Reintroducing Samba
 - ▶ Making a file server
 - ▶ Troubleshooting
-

There are some things you can't share without ending up liking each other.

— J. K. Rowling, *Harry Potter and the Sorcerer's Stone*, 1997

Besides printers, the next big thing in local networking is file sharing. File sharing is a necessity in the business world, but it can also come in handy within the home network as well. Having a common place for everyone in the house to place pictures, music, or documents makes using floppies and flash drives almost obsolete.

There's no doubt that the Microsoft network environment is the most popular file sharing system available. Any Microsoft Windows computers can connect together to share folders with one another. But just because you decided to use Linux doesn't mean you're left out in the cold. In this chapter, I walk you through the steps necessary to create your own file server on your home network that can be accessed from any Microsoft or Linux device on your network.

Reintroducing Samba

In Chapter 20, you saw how the Samba software package could be used to share a local printer with Microsoft Windows clients on the network. The same Samba software package also makes sharing folders and files between Linux and Windows devices a snap.

Samba allows you to define folders on your Linux system that it then advertises on the Microsoft network. Remote Windows clients can browse and map to those folders just as if the folders were on a Windows system.

You can control access to the folders by using the standard Linux user and group permissions. Samba has the ability to allow Windows users on the network to access individual Linux user accounts. This feature allows you to use one set of user accounts on the Windows network, and a different set of user accounts on the Linux system.

Step 1: Plan a Sharing Policy

Before you can start sharing, first you must decide what you want to share and who gets to share it. Mapping out a folder sharing policy is a crucial element to creating a file server.

Even in small organizations, you must often restrict subgroups from accessing certain information. Whether it's separating teacher data from students, or keeping children away from parents' data, you often find yourself creating more than one group for sharing folders.

You can create as many shares on the file server as you need to manage the data groups. For each group that must share data, you must create a separate folder and define a separate group to place users in.

With all this information to keep straight, it's usually best to create a simple table, mapping out what users should have access to what data. After you visually see the layout, it is easier to create the proper configuration. Table 21-1 demonstrates a simple shared folder policy table.

Table 21-1 Mapping Out a Shared Folder Policy			
Group	Folder	Share Name	Members
accounting	/accounting	accounting	fred, barney
sales	/sales	sales	wilma, betty

Table 21-1 shows that each group is assigned a unique folder on the Linux system. To keep things simple, this example uses the same group, folder, and share names, but on a real system, this may not be possible. Making a table that maps out what’s what can save you from lots of headaches later on in the project.

Step 2: Create Linux Groups

The next step in creating your file server is to create the necessary groups in Linux. The file server uses the Linux groups to control access to the individual shared folders that you’ll create. For this step, you need to use the Fedora User and Group Manager:

- 1. **Select System⇨Administration⇨Users and Groups from the panel menu.**

The User Manager window appears. If you are not logged in as the root user, Fedora first asks you for the root password to continue.

This tool provides a graphical interface for creating and managing users and groups on the system.

- 2. **Click the Add Group icon to add a new Linux group.**

The Create New Group window, shown in Figure 21-1 appears.



Figure 21-1:
The Fedora
Create New
Group
window.

In the text box, enter the name of the group you want to use to control the shared folder. Leave the Specify ID manually check box blank. Fedora automatically assigns a unique group number to the group.

Use this technique to create a group called accounting, and then another group called sales. The new groups appear in the group listing under the Groups tab.

Step 3: Create Folders to Share

After you've created the groups, you can create the folders to share. For security reasons, it's not a good idea to create folders under an existing user's home folder. Instead, you should create folders separate from the normal directory structure on the system.

- 1. Log in as the root user.**

This is one of the times where it is easier to just log in as the root user. Anytime you log in as root, you must use extreme caution that you don't alter anything you shouldn't be altering.

- 2. Double-click the Computer desktop icon, and then double-click the Filesystem icon.**

This starts Nautilus, showing the root directory structure of the Linux system.

- 3. Right-click on an empty space in the Nautilus window and select Create Folder.**

When Fedora creates the new folder, it gives you the option to name it. Change the name to accounting.

- 4. Create another folder, called sales, in the root directory.**

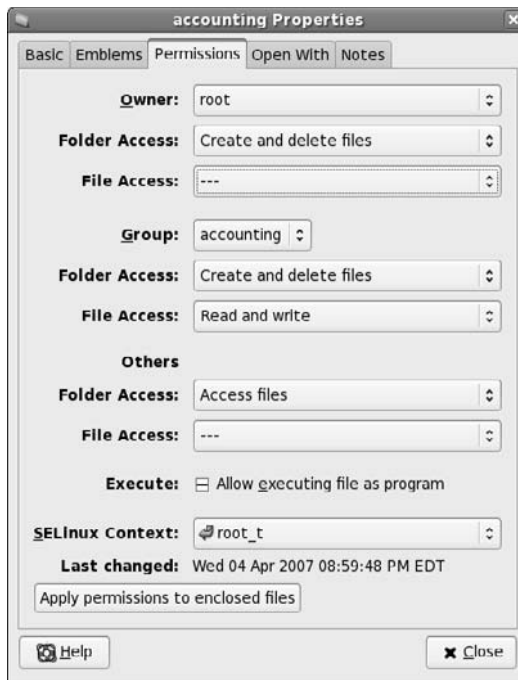
- 5. Right-click on the accounting folder and select Properties from the drop-down menu.**

The accounting Properties window displays the settings for the folder.

- 6. Click the Permissions tab.**

The Permissions tab displays the permissions for the owner of the folder, the assigned group the folder belongs to, and access for all other users on the system. This is shown in Figure 21-2.

Figure 21-2:
The
accounting
Properties
window
with the
Permissions
tab
selected.



7. Change the group setting to accounting.

By default, Fedora sets the folder's group to the group of the user who created it, which in this case is root. Use the drop-down box to change the value to the accounting group you previously created.

8. Select Create and Delete Files from the Folder Access drop-down list, and select Read and Write from the File Access drop-down list.

This allows anyone in the accounting group to create new files and remove existing files. Close this window when you are done.

9. Do the same for the sales folder, assigning it to the sales group, and giving the group Create and Delete Permissions.

Step 4: Set Advanced Folder Permissions

There is one additional trick that you must use to complete the shared folder permissions. When Fedora Linux creates a user account, it assigns each user to a unique group (which happens to be the same as the username). Unfortunately, this causes problems in a shared folder environment.

Even though your users are all members of the same group, their primary Linux group is still their unique group name. When a user creates a new file in the shared folder, it belongs to the user's primary group, not your newly created group.

To solve this problem, you must tell the shared folder to use the folder's group, not the user's group, when users create new files. This requires a special Linux permission that is not available from the Nautilus Properties window.

You must manually assign this permission from the command line prompt by using the Linux `chmod` command. Here are the steps to accomplish this:

1. **Select Applications→System Tools→Terminal to start the command line session.**
2. **If you are not logged in as the root user, from the command line, type `su` to become the root user. Enter the root user password when prompted.**
3. **Type the command `chmod g+rws /accounting`.**

The `chmod` Linux command changes the permissions on the specified folder. The individual options in the command are

- `g+` specifies for the permissions to be added to the group permission settings.
- `r` specifies for the group to have read permissions to the folder.
- `w` specifies for the group to have write permissions to the folder.
- `s` specifies for the folder group setting to be used for all files created within the folder.

4. **Use the `chmod` command again, this time using the `/sales` folder.**

Now whenever a user creates a file in the accounting or sales folders, the group of the file is set to the special group name instead of to the user's primary group.

This process ensures that other members of the group can access files created by one member of the group.

5. **Type `exit` at the command line prompt to exit root mode, and then type `exit` again to exit the Terminal session.**

Step 5: Install and Start Samba

If you haven't already installed and started Samba while working in Chapter 20, you'll need to do that now. You can use the Fedora Package Management tool to install the Samba software:

1. **Click Applications→Add/Remove Software.**

This starts the Fedora Package Management system. You can easily add or remove software packages from here. If you are not logged in as the root user, it will ask you for the root user password to continue.

2. **In the Package Management window, select the Servers option on the left side and the Windows File System option on the right side.**

The Windows File System option includes one optional package — system-config-samba. Ensure that you select this package for installation as well. It will make configuring the Samba software easier.

3. **Click the Apply button on the lower right to start the software installation.**

After you install the Samba software, you must decide how you want it to run on your server. You have two options for this:

- ✓ **Manually:** If you don't want the Samba server running all the time
- ✓ **Automatically:** If you want the Samba server to start when the Linux server is booted

For either option, you must use the Fedora Service Configuration window.

1. **Select System→Administration→Server Settings→Services from the menu.**


This starts the Service Configuration window. Again, you must be the root user to use this tool, so if you're not already logged in as root, you'll be asked for the root password to start. When the Service Configuration window starts, you see all the services that are available on the Fedora system, as shown in Figure 21-3.

The server automatically starts services that have the green check mark in the check box when the system boots. You can also start, stop, or restart individual services by using the appropriate button.

The Services Configuration window labels the Samba service as smb. This service actually controls two separate Samba services:

- **smbd:** Controls client connections to the Samba server.
- **nmdb:** Controls advertising shared devices on the network.

Figure 21-3:
The Fedora
Service
Configu-
ration
window.



2. Select the smb service, and click Start.

When you start the smb service, both of the Samba servers start. You will see the status of the services in the right window.

3. If you want Samba to automatically start at boot up, select the check box next to smb, and then exit the Service Configuration window.

Step 6: Configure Samba

There are two main features that you must configure to get your Samba server to work on your network:

- ✓ The Microsoft network domain or workgroup name
- ✓ Opening the Samba server firewall port to network clients

Here are the steps you need to follow to accomplish this.

1. Select System⇨Administration⇨Server Settings⇨Samba from the menu.

This starts the Samba Server Configuration window.

2. Click Preferences⇨Server settings from the menu bar.

The Server Settings configuration window, shown in Figure 21-4, appears.

Figure 21-4:
The Samba
Server
Settings
configu-
ration
window.



In the Server Settings window, enter the workgroup name for your Windows network (remember, all network devices that need to use the shared folder must use this same name).

By default, Fedora uses the Linux hostname of your system as the Windows network device name on the network. You can also provide a brief description of your host in the Description dialog box provided. Network devices see the description along with your hostname when they browse the network.

3. Click OK when you are done, and then close out the Samba Settings window.

The next step is to open the firewall port that Samba uses to allow remote clients to access the server. This is done by using the Fedora Security Level Configuration window:

4. Select System⇨Administration⇨Firewall & SELinux from the menu.

This starts the Security Level Configuration window, as shown in Figure 21-5.

Figure 21-5:
The Fedora
Security
Level
Configu-
ration
window.



You must be the root user to have access to this tool, so if you're not logged in, you must enter the root password. This tool has two tabs:

- **Firewall Options:** For opening specific firewall ports
- **SELinux:** For controlling how Linux locks specific applications

5. In the Firewall Options area, select the Samba check box.

The Firewall Options area shows a list of common services that are configured on the system. Selecting the Samba entry marks it as a trusted service and allows network clients to connect to it.



If you enabled the SELinux security software during installation, you'll need to tell it to allow Samba to operate. To do that, select System⇨Administration⇨SELinux Management to start the SELinux Management tool. From there, select Boolean, and then click the Samba drop-down list box. Select the Allow Samba to Share Any File/Directory Read/Write check box. That's all there is to it!

Step 7: Create UNIX User Accounts

Now that your Samba server is live on the network, it's time to turn your attention back to setting up the shared folder environment. You've already created the groups necessary for access to the shared folders, so now it's time to work on the user accounts.

Each user who needs to access a shared folder must have an individual Linux user account. As you saw in Chapter 17, you can do this with the User Manager:

1. Select System⇨Administration⇨Users and Groups from the panel menu.

This starts the User Manager window. If you are not the root user, it first asks you for the root user password.

2. Click the Add User button to create a new user.

The Create New User window, shown in Figure 21-6, appears.

Enter a username for the new user, along with the user's full name for documentation purposes. You have to assign a password to the user account, but don't worry: Samba won't use this password to authenticate the remote client.

3. Add new user accounts for each user who needs access to the shared folders.

To create the users defined in the sample share plan, create users fred, wilma, barney, and betty.

4. Click the Groups tab to view the configured groups.

The Groups tab displays the groups that are configured on the server. You should see the groups you created back in Step 2.

5. Add the new users to the appropriate groups.

Double-click the accounting group name. The Group Properties window appears, allowing you to select the users that belong to the group under the Group Users tab. Select the users who need to belong to the accounting group (fred and barney), and then click OK. Do the same to add wilma and betty to the sales group.



Figure 21-6:
The Fedora
Create New
User
window.

Step 8: Create Samba User Accounts

I mentioned in the preceding section that Samba doesn't use the Linux user account password when authenticating remote clients. That's because Samba uses its own user accounts for that. Samba acts like a middleman between the remote Windows client and the Linux filesystem.

When accessing files and folders on the system, Samba uses the defined Linux user accounts. When communicating with remote clients, Samba uses its own internal user accounts. Because of this, Samba must map the Windows user accounts to the appropriate local Linux user accounts. This is done in the Samba Server Configuration window:

1. Select **System**⇨**Administration**⇨**Server Settings**⇨**Samba** to start the **Samba Server Configuration** window.

As before, you need root permission to use this program, so if you're not already logged in as the root user, Fedora will ask you for the root password.

2. Click **Preferences** → **Samba Users** from the menu bar.

This displays the Samba Users window, as shown in Figure 21-7.

Figure 21-7:
The Samba
Users
window for
adding new
Samba
users.



3. Click the **Add User** button to add a new Samba user.

The Create New Samba User window appears. This window contains four areas where you need to enter information:

- **UNIX Username:** A drop-down list box allowing you to select the Linux user account to map.
- **Windows Username:** A text box to enter the associated Windows network username the remote client will use.
- **Samba Password:** The password assigned to the Windows network username (remember, this can be, and often is, different than the Linux password).
- **Confirm Samba Password:** Retype the Samba Password to ensure that you typed it correctly.

Remember, you must match the Windows network username with the Linux username. In the example shown in Figure 21-8, Samba maps the Windows network username fflinsto to the Linux user account fred.

Figure 21-8:
Creating
a new
Samba user
account.



4. (Optional.) If you need to modify a Samba account, select the username from the Samba Users list and click the Edit button.

After mapping the accounts, you can leave the Samba Server Configuration window open because you need it for the next step.

Step 9: Create Shared Folders in Samba

The final piece of the puzzle is to tell Samba what shared folders we want to advertise on the network and allow network clients to connect to. This is done in the Samba Configuration window.

1. Click the Add Share button.

The Create Samba Share window appears. From here, you define the details of the shared folder:

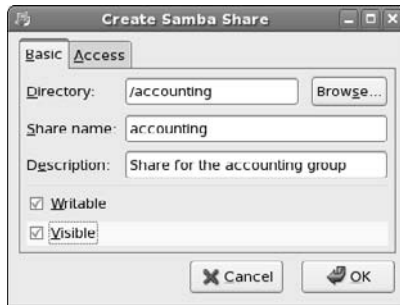
- **Directory:** Type the full pathname of the folder you want to share. You can also click the Browse button and browse to select the folder.
- **Share Name:** Type the name you want Windows network clients to use to connect to the share.
- **Description:** Type a simple description that will be advertised along with the share name to remote clients.
- **Writable:** Check this check box if you want to allow remote clients to be able to write to the shared folder. You must also have the Linux permissions set, as shown in the earlier sections in Steps 3 and 4, for them to write. This is a handy way to stop allowing remote clients to write to the share in case of an emergency.
- **Visible:** Select this check box if you want remote clients to see the share advertised when they browse the network.

Enter the /accounting folder for the directory. As you type the folder name, Samba enters a suggested share name as well (which is usually the same as the folder name). Figure 21-9 demonstrates configuring the accounting shared folder.

2. Click the Access tab and select which users have access to the shared folder.

Samba allows you to assign access for individual Linux user accounts configured on the computer. You can either select the check box to allow access for all users, or select individual users from the list. Allowing all users isn't a bad thing because the Linux permissions still block out unauthorized users from the share.

Figure 21-9:
Sharing a
folder in
Samba.



Step 10: Map to Shared Folders

Now that your file server is complete, it's time to test it out. There are a few different ways you can do this. The easiest is to test mapping a shared folder from a remote Windows client. For this test, I'm using a Windows XP Home Edition client:

1. **Open the My Computer icon by either double-clicking the desktop My Computer icon or selecting it from the Start menu.**
2. **Click Tools⇨Map Network Drive from the My Computer menu bar.**
3. **Click the Browse button to browse the network.**

Under Microsoft Windows networks, you should see the Fedora Linux server listed as a network resource. Expand the resource and you'll see the share names listed. Select the share you have access to.

4. **Select an available local drive letter to map the network shared folder to.**
5. **Click Finish to complete the network shared folder mapping.**

That's all there is to it. You should now be able to see the shared folder in your My Computer area. Clicking the shared drive letter shows the contents of the shared folder. You should have the ability to create, modify, and delete any files in that folder.

Troubleshooting

As with everything else with computers, there are always places for things to go wrong while creating a file server. Here's a brief list of some places to watch when things don't work as planned:



- ✓ **Users can't access their shared folders:** Check all the permissions. First, check to make sure that the proper Linux group is set for the folder, and then check to make sure that the group has the proper permissions for the folder. Finally, check to make sure that the individual users are members of the Linux group.

After you've verified that the Linux permissions are OK, look into the Samba permissions. Check to make sure that the proper Samba user accounts have been given access to the share.

- ✓ **Users can't see shares when they browse the network:** Check to make sure that the Visible check box is enabled for the share in the Samba Configuration window. Selecting this check box allows Samba to advertise the share on the network. If this check box is not selected, the share will not be advertised, and it won't appear when clients browse the network.

This can be used as an extra security feature. Hackers can't try to break into what they can't see. Valid users can map to the share manually by using the Microsoft network name `\\servername\sharename`.

- ✓ **Users can't see the file server when they browse the network:** Make sure that the smb service is enabled in the Services Configuration window. If the smb service is not enabled, remote clients won't see the server on the network and won't be able to connect to it.
- ✓ **Users can see the file server, but can't map to any shares:** Make sure the Fedora firewall settings have Samba enabled. If the Samba port is not enabled, Samba can send packets, but not receive them.

Chapter 22

Ten Troubleshooting Tips

In This Chapter

- ▶ Dealing with frozen installations
 - ▶ Changing your boot environment
 - ▶ Recovering when you see some black-and-white text screen!
 - ▶ Escaping a hung GUI
 - ▶ Using multiple resolutions
-

Troubleshooting is like reading a mystery novel. You have some facts, symptoms, and details, but you don't know whodunit. You have to take whatever information you have, work with that data, weigh the various possibilities, and then narrow them to a single suspect. Finally, you need to test your theory and prove that your suspect is the guilty party.

Troubleshooting problems in Linux (or any operating system) can encompass many hardware and software issues. Whether the problem is the operating system, the hardware, or a service giving you fits, you can use some basic troubleshooting techniques to start your investigations:

- ✓ **Document the problem.** Write down any and all symptoms that the system is showing, including actions you can and can't do. Jot down any information you see in error messages.
- ✓ **Examine the Linux log files.** You can find most of these in the `/var/log` directory. Look for the word "error."
- ✓ **Compare your problem system with a working system running the same distribution and version.** Sometimes, comparing configuration files and settings may uncover the problem or narrow the possibilities.
- ✓ **Check connections.** Check to make sure that all the hardware is connected properly and powered on. Verify that all cables and connections are attached properly. There's always someone, somewhere, accidentally kicking a cable out from a wall connection.

- ✓ **Remove new hardware.** Remove any hardware that you have changed or added recently (before the problem started) and see whether the problem disappears. If so, you can probably conclude that the new or changed hardware (or its driver) is the culprit and start researching solutions.
- ✓ **Reduce the number of active programs.** Stop running unnecessary services and applications that aren't related to the problem at hand. You may more easily figure out what's happening if other services and applications aren't getting in the way.
- ✓ **Check to see whether the problem is reproducible.** Does the same sequence of events produce the same problem? Suppose that when you try to print to a color printer, nothing happens. If nothing happens *every* time you attempt to print, the problem is reproducible. If, instead, sometimes your information is printed and at other times it isn't, the problem pattern isn't the same and isn't reproducible — or it's caused by something more complicated than just clicking one button. Unfortunately, problems that are nonreproducible are more difficult to resolve because it seems that no set pattern of events re-creates those problems.

After you've come up with a solution, take a few moments to document the situation. Note the symptoms of the problem, its cause, and the solution you implement. The next time you encounter the same problem, you can call on your notes for a solution rather than reinvent the wheel.



If you don't have any problems to troubleshoot (yet), document your environment *before* you do. Making a backup of your `/etc` directory and your `/boot` directory is a great place to start.

Tip #1: “The Linux Installer Froze”

When you're installing Linux, the installation may just freeze. If it does, wait a bit and make sure that the installation program really froze. (Sometimes, the software just takes a while to process information.) If the software looks like it has frozen, there's no harm in rebooting your computer and starting over — just as you would do with any operating system installation. Sometimes, you can reboot and never have that problem again. At other times, the problem may happen twice in a row and then be fine the third time. Be sure to try several times before giving up.

If the installation still freezes in the same spot or close to the same spot, go to the distribution's support pages (see Chapter 2). These pages may talk about some known problems and solutions that can help you and should

show you how to join discussion lists in order to get more assistance. Otherwise, diagnosing the problem can be tricky and may seem more like voodoo than science. Here are some tips:

- ✓ **If this problem happens repeatedly at exactly the same spot, you may have a bad installation disk.** See the next section, “Checking Your Distribution Burns,” and then return here if that technique doesn’t solve your problem. Otherwise, try the disk in another machine if possible and see whether the installation fails in the same place there. If you purchased this disk from someone, contact their technical support team. If you got the disk with a book, contact the publisher’s technical support team. If you burned the disk yourself, try burning a new copy at a slower speed.
- ✓ **If this problem happens repeatedly at exactly the same spot and you don’t have a bad installation disk, the trouble may be with one of your machine’s hardware components.** If you can, try trading hardware between machines. If not, you may need to choose a different machine on which to install Linux or try another distribution.
- ✓ **If the problem seems to happen randomly, your particular Linux distribution may not be compatible with that particular machine.** Again, you can try trading some hardware around, installing Linux on another machine, or using another distribution.



If you’re not sure whether your installer has frozen, try pressing various combinations of Alt+F#, where # corresponds to one of the function keys. Depending on the distribution, the installer has not completely frozen if you can see different screens when you try this technique.

Tip #2: Checking Your Distribution Burns

When installing Red Hat Linux 7.3 or later (all versions of Fedora Core are considered “later”), a special solution is available to people who run into problems that seem to have absolutely no explanation, such as the installer freezing.

If your installation keeps dying while Anaconda (the Red Hat installer program) is placing packages on your hard drive, follow these steps to try to fix it:

1. **Place the DVD-ROM or the first Fedora CD-ROM into your drive.**
2. **Reboot the machine.**

3. Wait until you reach the Media Check dialog box.

If you've changed your mind and just want to start the installation, use the Tab or arrow keys to select Skip and then press Enter.

4. Tab to the Test button and press Enter. If you are using CD-ROMs and want to test a different one, tab to the Eject CD button and press Enter.

If you chose the second option, remove the first CD-ROM from the CD-ROM drive and replace it with the CD-ROM you want to test. Close the CD-ROM drive and make sure that Test is selected.

5. Press Enter to begin the media check.

The Media Check status box opens and shows you the name assigned to the DVD-ROM or CD-ROM and how much progress has been made. At the end of the inspection, the Media Check Result dialog box opens.

6. Look at the text after and the result is.

If the result is `PASS`, nothing is wrong with the DVD-ROM or CD-ROM itself. Your installation woes are caused by something else. Return to the section "The Linux Installer Froze," earlier in this chapter.

If the result is `FAIL`, the DVD-ROM or CD-ROM you just tested is flawed. If you purchased this CD-ROM or DVD-ROM, you need to talk to the company you purchased it from to see whether you can get a replacement. On the other hand, if you burned your own DVD-ROM or CD-ROM, I recommend doing one of the following:

- Burn the DVD-ROM or CD-ROM again, at a speed of 4x or lower.
- Burn the DVD-ROM or CD-ROM again on a newer drive with BurnProof technology (www.digital-sanyo.com/BURN-Proof) or something similar.



If the DVD-ROM that came with this book is defective, contact the technical support address listed in this book, not Red Hat. However, remember that the media-checking routine does not give you a reliable `PASS` or `FAIL` for this particular DVD-ROM.

Tip #3: "I Told the Installer to Test My Graphics, and They Failed"

The installer may have guessed wrong about what hardware you have. Double-check the settings as best you can. If they look right, try choosing a

lower resolution for now and testing again, and if that fails try a lower number of colors and test again. You can then try setting things back the way you want them after the machine is fully installed and updated, when it hopefully will have a fix for whatever the problem might be.

Tip #4: “The Installer Tested My Graphics Fine, but My GUI Won’t Start”

If your Linux installation program showed you a GUI test screen saying that you were ready to proceed with the rest of the installation, you probably expected that the GUI would start with no problem. Unfortunately, that doesn’t always happen.

If you boot your machine for the first time and see error messages when you’re trying to enter the GUI automatically or when you type `startx` to start the GUI manually, type `system-config-display` at a command prompt to start a program that will help you fix the problem.

Tip #5: “I Think I’m in Linux, but I Don’t Know What to Do!”

Two different screens tend to cause panic to folks new to Linux. The first of these screens, shown in Figure 22-1, is in fact a sign that you installed the software and booted the machine successfully. Jump for joy! It’s just that you’re booting into the command-line environment rather than the GUI environment. If you reach a screen similar to the one shown in Figure 22-1, the computer is asking you to log in with the username for an account and a password that you created during the installation process.



If you created only the root account, you can log in there as `root`.

After you enter the username and password, you find yourself at the screen shown in Figure 22-2, which just happens to be the second spot where people get worried. If you see this screen, you have not only booted properly into Linux, but you’re also logged in and using the machine! Give yourself a good pat on the back.

Figure 22-1:
A Linux
command-
line login
prompt.

```
[dee@localhost ~]$ telnet localhost
Trying 127.0.0.1...
Connected to localhost.
Escape character is '^'.
Fedora release 6.93 (Rawhide)
Kernel 2.6.21-1.3116.fc7 on an i686
login:
```

Figure 22-2:
Logged in at
the Linux
command
line.

```
[dee@localhost ~]$ telnet localhost
Trying 127.0.0.1...
Connected to localhost.
Escape character is '^'.
Fedora release 6.93 (Rawhide)
Kernel 2.6.21-1.3116.fc7 on an i686
login: dee
Password:
[dee@localhost ~]$
```

What do you do from here? Anything you want. Surf through this book for commands you want to run. Type `startx` to start up the GUI. If you didn't install any GUI (which means you selected a minimal install option with no graphical interface, or you actually unselected graphics), you may want to reinstall, or you have to add all the tools by hand (which is not a quick job!).

Tip #6: "I Don't Want to Boot into This!"

Are you booting into the command-line environment when you want to use only the GUI? Or are you finding that you're already booting into the GUI and you would rather boot to that nice, clean, black-and-white command-line screen? You're not stuck with either of these options. You can change them at any time.



You can press Ctrl+Alt+F# (F# refers to function keys F2 through F6) to change out of the GUI to a command line terminal at any time and then Alt+F7 or Alt+F8 to switch back.

Tip #7: Changing Your Boot Environment “Permanently”

The word *permanently* is in quotes in the heading because you can, of course, go back and change this setting later, if you want. *Permanently* just refers to the fact that after you have made this change, every time you boot the system, it automatically goes into the preferred environment until you change it.



You can’t make this change in Linspire, Ubuntu, or Xandros unless you want to boot into *single-user mode*, which is basically “safe mode” and not of much use. Well, okay, you can make the change, but you would need a techie friend to set up a bunch of stuff for you. Linspire and Xandros assume that you don’t want to do this. You can’t do this in Knoppix either, but then it’s a LiveCD so you won’t be doing anything permanent to it.

To make this change in Fedora, Mandriva, or SuSE, you need to edit what’s called a *run level*. Fortunately, all three of these distributions use the same run level settings, so the instructions are the same for all of them:

- 1. In the GUI, open a command line terminal.**

If you’re not sure how to do so, see Chapter 6. If you’re not in the GUI and you’re already logged in, type `su -` to become the root user.

- 2. Type `cp/etc/inittab/etc/inittab.old` to make a backup.**

Now, if something happens while you’re editing the `inittab` file, you can always restart fresh with the old version.

- 3. Open the `inittab` file in your preferred text editor.**

Some Linux text editors are covered in Chapter 11.

- 4. Scroll down until you find a line similar to the following:**

```
id:5:initdefault:
```

This line appears near the top of the file. What you’re interested in here is the number. In most mainstream Linux distributions, the number 5 tells Linux to boot into the GUI, and the number 3 tells Linux to boot

into the command line. In the preceding example, therefore, I boot into the GUI.

5. Change the number in this line.

If it's a 5, change it to 3, and vice versa. Make sure that all colons and other items are left properly in place, or else your machine will have problems booting later.

6. Save and exit the file.

The changes go into effect the next time you reboot the system.



If you do end up having problems booting the system, Fedora's DVD or first CD can be used as an emergency boot disk.

Changing your boot environment just for now

At any time, you can have your Linux box switch between full command-line mode and full GUI mode.

To switch between modes, do the following:

- ✓ To change from the GUI login to the command-line login, open a terminal window and type (as root) `init 3`.
- ✓ To change from the command line login to the GUI login, type (as root) `init 5`.

Tip #8: "I Want to Change Screen Resolutions"

Do you want or need to swap between resolutions in the GUI on the fly? Suppose that you want to use 1,024 x 768, but you work on Web pages and want to be able to see how they look in a browser at 800 x 600 or even 640 x 480. Your machine is very likely already set up to do this, but you just need to know how!

If your machine is set up for it, you can change resolutions by pressing the key combination `Ctrl+Alt+Plus`, where Plus is the big plus (+) sign on your number pad — you can't use the plus sign on the main keyboard for this one.

If you're using a keyboard without a number pad — as you will be if you're using a laptop — or your machine isn't set up to be able to change on the fly, you need to change your resolution through your display configuration program. You can find this option at System⇨Administration⇨Display.

Tip #9: "My GUI Is Hung, and I'm Stuck!"

One quick solution to this problem is the key combination Ctrl+Alt+Backspace. If this doesn't do the trick, your system is in really bad shape! Try to switch to a virtual terminal by using Ctrl+Alt+F5. If this key combination also does nothing, you need to reboot the machine.

Tip #10: "Help, My Machine Hangs During Boot!"

When configuring a Linux machine, you may encounter problems with the `/etc/grub.conf` file. This file indicates the operating system or systems to which your system can boot, and the file also contains Linux start-up settings. In order to fix your computer, you can either try rebooting and selecting a different Linux boot option from the menu, or you could refer to Chapter 4 for instructions on using the rescue disk to boot into rescue mode. Consider this list of potential solutions if the `/etc/grub.conf` file makes trouble:

- ✓ If you have altered or added hard drives, you may need to change the boot line in the `/etc/grub.conf` file.
- ✓ If you haven't made hardware changes, check to make sure that your `/etc/grub.conf` file is referring to the correct location of the Linux image. (The program code that loads and executes at runtime and is located in the `/boot` directory.)
- ✓ If the location under the `/boot` directory or the device for the root entry is incorrect, your system can't boot to Linux.
- ✓ If you're working with a multiboot operating system environment, be sure that your `/etc/grub.conf` file contains entries for each of your operating systems. Each operating system or Linux installation needs to be in separate entries.
- ✓ If your file contains entries to switch to a higher-resolution display and you have boot problems, try reducing the video setting to simple VGA.



Linux allows you to use spaces and other characters in filenames that you may or may not be able to use in filenames on other operating systems. However, some Linux applications may stumble when they encounter file or directory names containing spaces. Usually, a safe bet is to stick with alphanumeric characters and avoid spaces and odd characters, such as question marks and exclamation points.

“Aaargh! I Forgot My Root Password! What Do I Do?”

Fear not. You have a way around this problem! You need to boot into *single-user mode*, which you can accomplish by rebooting your machine. When you see the blue screen with the words “Press any key to enter the menu,” press a key. At the GRUB boot screen, press **E**, which takes you to a configuration file. Use the arrow keys to go to the line starting with `kernel`, and press **E** again to edit that line. At the end of the line, add the word `single`, press Enter to put the change into place, and then press **B** to boot the machine.

Type `passwd` and then enter the new password twice as directed. When you finish, type `exit` and then boot the machine normally.

Appendix A

Common Linux Commands

In This Appendix

- ▶ Archiving and compression commands
 - ▶ Communication commands
 - ▶ File commands
 - ▶ System control commands
-

Computing novices often marvel at the keyboard dance Linux experts typically perform. Sure, these experts know about modern advances like the mouse and graphical interface, but these keyboard musicians prefer the home keys and find that they can work faster that way. It takes some time to reach this level of proficiency, but every expert was a novice at one time, and any novice can become an expert with plenty of practice.

In this appendix, you find the commands listed by themes, according to what they can actually do for you.

So, read on — and dazzle your friends with your command-prompt finesse. When they ask you how and where you figured out all those commands, just smile and mumble something about the voices in your head — and, of course, keep this section dog-eared and within reach of your computer.

Linux Commands by Function

Because every command serves a specific purpose, organizing these tools into groups according to their individual functions isn't difficult. If you know what you need to do but don't know which command does the job, flip through this section to start your search. From here, you can dig further by referencing `man` pages and other help information (online sites and reference books, for example) or by looking in this book's index for further coverage.

To access a man page, type **man *command*** at a command prompt. For example, `man ls` shows you the help information for the file listing command.

Getting Help

When you're digging around for help on a command, you can call on an interesting range of shell commands for assistance, as shown in Table A-1.

Table A-1 Shell Help Commands	
Command	Purpose
<code>apropos</code>	Looks for commands that contain a keyword in their man page descriptions.
<code>info</code>	One way of finding help information. You can find instructions for this tool at www.gnu.org/software/texinfo/manual/info/ , or you can use the built-in tutorial by starting the <code>info</code> tool and pressing the H key when inside it.
<code>man</code>	The primary way of getting help in Linux and UNIX.
<code>whatis</code>	Gets a one-line description of a command.

Locating details about the command-prompt options of a command is a never-ending pursuit. The man page system provides some helpful guides at your fingertips for rapidly finding this detailed information.

Archiving and compressing

Although disk space isn't as much of a premium as it once was, bandwidth and backup media still are. Subsequently, this group provides a potpourri of tools for compacting and organizing data for storage, as shown in Table A-2.

Table A-2 Archiving and Compressing Tools	
Command	Purpose
<code>bzip2</code>	Compresses files into <code>.bz2</code> format. Used mostly for incredibly large sets of text files (which is what source code actually is).
<code>bunzip2</code>	Uncompresses <code>.bz2</code> files.

Command	Purpose
<code>compress</code>	Compresses files into <code>.z</code> format. Pretty old and not used much in the Linux world.
<code>gunzip</code>	Uncompresses <code>.gz</code> files and <code>.tgz</code> files.
<code>gzip</code>	Compresses files into <code>.gz</code> format.
<code>tar</code>	Packages files together in a group. The most common way of using this command is <code>tar xvf filename</code> , such as <code>tar xvf download.tar</code> .
<code>uncompress</code>	Uncompresses files from <code>.Z</code> format.
<code>unzip</code>	Uncompresses files from <code>.zip</code> format.
<code>zip</code>	Compresses files into <code>.zip</code> format.

Built-in bash commands

Some commands don't even seem to exist if you try to look up their help information in the `man` pages, and the commands don't show up as files on your system. Remember, as you type commands at the prompt, that you're communicating with a type of program called a *shell*. (In my case, it's `bash`, the default Linux shell.) The shell has a set of commands, included in the following list, that you can use to communicate with it, as shown in Table A-3.

Table A-3 Shell Commands	
Command	Purpose
<code>alias</code>	Creates or lists command shortcuts.
<code>env</code>	Lists your current environment variables and their settings.
<code>export</code>	Whenever you're told to set an environment variable, create the variable and then use this command so that the variable will be remembered properly.
<code>history</code>	Lists off the last 1,000 commands you've typed.
<code>unalias</code>	Removes command shortcuts.



If you try to view the `man` page entry for some of these commands, you find instead the help information for `BASH BUILTINS` loads. To search through this manual, press the forward slash (`/`) key to open the `man` search interface and then type the name of the command you want to search for. Press `Enter` to start the search. The interface stops in the first spot where the term is found. If you want to try again, press the `N` key to proceed to the next occurrence of the word.

For example, you might be reading the massive `bash` `man` page (type `man bash` to access this page), but perhaps you’re only interested in items related to *prompts*, which are the bits of text that appear to the left of your cursor in a text window. An example prompt is

```
[dee@catherine dee]$
```

So, you might type `/prompt` and press `Enter` to jump down to the first instance of this word. If the text around the word doesn’t reflect what you’re looking for, you press the `N` key to jump to the next one, and so on.

Files and Filesystem

No matter which operating system you’re using, it’s hard to do anything without being able to find your way through and work with the filesystem. The following utilities help you find your way:

File organization

Boxing, packing, sorting, shipping — I’m always shuffling files around on my system. File organization commands provide tools for moving files and filesystem units around, as shown in Table A-4.

Table A-4 File Organization Tools	
Command	Purpose
<code>cd</code>	Changes directories.
<code>cp</code>	Copies a file.
<code>df</code>	Shows partitions and how much space they have.
<code>du</code>	Shows how much disk is being used in the current directory and below.

<i>Command</i>	<i>Purpose</i>
<code>ln</code>	Creates a shortcut.
<code>ls</code>	Lists the contents of a directory or information about a file.
<code>mkdir</code>	Creates a directory.
<code>mv</code>	Moves or renames a file.
<code>pwd</code>	Shows the path for the directory you're currently in.
<code>rm</code>	Deletes a file.
<code>rmdir</code>	Deletes an empty directory.

File attributes

Files are much like candy bars. The wrappers provide information about the ingredients, size, and package date — all descriptive of the tasty nugget inside. (Perhaps the wrapper is even childproof.) Files keep all this wrapper information in an *inode*. Along with the capability to change file inode information, these commands can return data about the content of the file, as shown in Table A-5.

Table A-5 File Attributes Commands	
<i>Command</i>	<i>Purpose</i>
<code>chgrp</code>	Changes the group associated with a file
<code>chmod</code>	Changes a file's permissions
<code>chown</code>	Changes who owns a file
<code>file</code>	Shows what type of file you're dealing with
<code>stat</code>	Shows some statistics about the file
<code>touch</code>	Creates an empty file of this name
<code>wc</code>	Shows how many words, lines, and so on are in this file

File locators

Where, oh, where can my file be? These commands, shown in Table A-6, help you locate files in Linux's monster tree-structure filesystem:

Table A-6 File Locators Commands	
Command	Purpose
find	Hard-core filesystem search tool.
locate	Lighter-weight filesystem search tool.
which	Tells you the path for the program that would be run if you typed this command.

File viewers

Text file browsing is a favorite pastime of many a system user. These tools provide a variety of utilities for viewing the contents of readable text files of all sizes. Unlike using a full-screen editor, you cannot damage the contents of a file with these commands, shown in Table A-7, because they’re just viewers, not editors:

Table A-7 File Viewers	
Command	Purpose
cat	Dumps the contents of the file to your screen.
head	Shows the first ten lines of a file.
less	Shows the file a screen at a time.
more	Shows the file a screen at a time.
tail	Shows the last ten lines of a file.

File-system commands

These commands, listed in Table A-8, provide information or perform actions on the entire filesystem, from creation and tuning to repair and recovery. Some of these commands return data only, whereas others also provide you with surgical instruments for serious file-system hacking:

Table A-8 File-System Commands	
<i>Command</i>	<i>Purpose</i>
<code>badblocks</code>	Searches a partition for bad blocks.
<code>e2fsck</code>	Checks and repairs an ext2 or ext3 filesystem.
<code>e2label</code>	Applies a filesystem label to an ext2 or ext3 partition.
<code>eject</code>	Ejects a CD or DVD.
<code>fsck</code>	Can check and repair many types of filesystems.
<code>mkfs</code>	Creates a filesystem (format a partition).
<code>mount</code>	Loads a partition into your filesystem.
<code>sync</code>	Saves all information out of buffers onto disks.
<code>tune2fs</code>	Adjusts ext2 and ext3 filesystem parameters.
<code>umount</code>	Removes a partition from the filesystem.

mtools

The mtools suite of utilities provides a nice way to transfer information to your Microsoft friends. Although Linux has native support for Microsoft Windows and DOS filesystems, your Microsoft cohorts don't have access to Linux (ext2 and ext3) filesystems. To keep everyone happy, you can buy preformatted MS-DOS disks and use them with the mtools commands (see Table A-9) so that you can swap them back and forth with your friends who are using Windows:

Table A-9 mTools Commands	
<i>Command</i>	<i>Purpose</i>
<code>mcd</code>	Changes directory in DOS format on a DOS disk.
<code>mcopy</code>	Copies DOS files to and from Linux.
<code>mdel</code>	Deletes a DOS file.
<code>mdeltree</code>	Deletes a DOS directory and its contents.
<code>mdir</code>	Lists a DOS directory's contents.
<code>mdu</code>	Shows how much space is taken and available for a DOS partition.

(continued)

Table A-9 (continued)	
Command	Purpose
mformat	Formats a partition for DOS.
mlabel	Applies a DOS volume label.
mkdir	Creates a DOS directory.
mmount	Mounts a DOS disk or partition.
mmove	Moves or renames a DOS file or directory.

System Control

These commands provide systemwide information and control. Normal users can run many commands to obtain system information; however, commands that actively change the configuration of the system need to run while you’re logged in as root — or have utilized the `su` command to temporarily become the superuser.

Administration

Some administration commands, shown in Table A-10, don’t fall neatly into a category.

Table A-10 Administration Commands	
Command	Purpose
passwd	Change a particular user’s password. Any user can run this command to change their own password. Only root can use it to change someone else’s.
su	Switch to another user account without logging out of this one. The best way to use this command is <code>su -</code> so your filesystem path and other information is loaded.

Kernel module handling

You may sometimes need to add kernel support for an additional device (software or hardware). If this need arises, you have a limited number of choices: You can either rebuild the kernel or install a loadable kernel module. Although rebuilding a kernel doesn't exactly require a PhD in nuclear science, consider it a time-consuming nuisance that's best to avoid. The commands in Table A-11 enable you to include the kernel support you need while the system is running, without having to rebuild the entire thing from scratch.

Table A-11 Kernel Support Commands	
<i>Command</i>	<i>Purpose</i>
<code>depmod</code>	Regenerates your module dependencies.
<code>insmod</code>	Loads a module by hand.
<code>lsmod</code>	Lists the modules your kernel has loaded.
<code>modprobe</code>	Loads a module by hand along with its dependencies and settings.
<code>rmmod</code>	Unloads a module by hand.

Processes

Most of your system activity requires processes. Even when your system appears idle, a dozen or so processes (programs) are running in the background. These commands, shown in Table A-12, enable you to check under the hood to make sure that everything that needs to be running is running and that you're not overheating or overtaxing resources:

Table A-12 Process Commands	
<i>Command</i>	<i>Purpose</i>
<code>crontab</code>	Sets up commands to run at regular intervals.
<code>kill</code>	Stops a process by its number. Often used as <code>kill -9</code> for a harsh stop for something that won't die.
<code>killall</code>	Stops a process by name rather than number.

(continued)

Table A-12 (continued)

<i>Command</i>	<i>Purpose</i>
nice	Assigns a CPU use priority to a process.
pidof	Gets a program's ID number.
ps	Gets a lot of programs' ID numbers, usually used as <code>ps aux</code> .
pstree	Shows the relationships between programs.
renice	Changes a program's CPU use priority.
top	Shows resource use over time.

Index

• Symbols and Numerics •

- * (asterisk) as wildcard, 88
- \ (backslash) in Windows versus Linux, 100
- (dash)
 - in file and directory names, 100
 - in permission sets, 108–110
- \$ (dollar sign)
 - in bash prompt, 84–85
 - retrieving variable values, 93
- . (dot)
 - in file and directory names, 100
 - starting hidden files, 86, 106
- " (double quotes) for strings, 94
- = (equals sign) for assigning variable values, 92
- / (forward slash) for directories, 100
- > (greater than sign) as redirection operator, 96
- | (pipe symbol) for redirecting output, 97
- ? (question mark) as wildcard, 88
- ' (single quotes) for strings, 94
- ~ (tilde) for home directory, 86
- _ (underscore) in file and directory names, 100
- 3-D games
 - Cedega for, 259–261
 - Linux challenges for, 253–254
 - Linux video card drivers for, 254–255
 - online, 255–256
 - repositories, 255
 - Second Life, 256

• A •

- access control for printers, 293
- administration. *See* system administration
- administrivia, 1

- Amarok music player
 - adding music files, 233
 - fetching covers for albums, 232
 - filtering collections, 232
 - finding and installing, 230–231
 - First-Run Wizard, 231
 - installing, 275
 - for Internet radio, 233–234
 - loading song lyrics, 231–232
 - playing music, 233
 - playlists, 232–233
 - tabs along window, 231–233
- Apache Web Server (on the DVD), 387
- apmd service, turning off, 304
- Apple QuickTime, 241
- applets (GNOME), 73–75
- applications. *See* programs
- ApplixWare Office suite, 192
- archiving commands, 378–379
- arguments with commands, 87
- asterisk (*) as wildcard, 88
- ATI binary drivers, 254
- attributes of files, commands for, 381
- audio. *See* music CDs; sound
- Audio Configuration tool, 227–228
- autocompletion feature, 66, 89–90
- AVI (Audio Visual Interleave) format, 240–241, 243

• B •

- b in `ls` listings, 110
- backing up
 - /etc/profile file before experimenting, 95
 - before partitioning, 29–30
- backslash (\) in Windows versus Linux, 100
- Barkakati, Naba (*Linux All-in-One Desk Reference For Dummies*), 1
- base directories, 100–102

Base (OpenOffice.org)
 creating a new database file, 213–214
 creating a table, 217–218
 Database pane, 217
 Database Wizard, 213–215
 described, 212
 menu bar, 215–216
 opening a database file, 214–215
 standard toolbar, 216
 Tables/Queries/Forms/Reports pane, 217
 Tasks pane, 217

bash (Bourne again shell). *See also* CLI
 (command line interface); commands
 autocompletion feature, 89–90
 built-in commands, 379–380
 clearing the screen, 98
 customizing the prompt, 93–94
 help commands, 378
 logging in, 85
 as most common shell, 83–84
 prompt for, 84–85
 resetting the prompt, 98
 starting a session, 84–85
 starting programs from, 87–88

Beagle search tool, 131–133

benevolent dictator, 15

BIOS (Basic Input/Output System), 38, 59

BitTorrent, 176–178

block devices in `ls` listings, 110

Blum, Richard (*Linux For Dummies*), 1–6

boot loaders
 described, 59
 options when installing Fedora, 47
 removing, 64
 troubleshooting, 60

boot menu, 21, 60

`/boot` partition setup, 47

booting. *See also* dual booting
 boot-loader problems, 60
 changing boot environment, 371–372
 as complex process, 58
 from DVD-ROM with this book, 389–390
 entering rescue mode, 62–63
 failed POST process symptoms, 58–59
 Fedora first time, 53–56
 getting BIOS information upon, 38
 getting video information upon, 38
 hang during, 373–374
 rebooting, 63–64

reinstalling Linux for problems with, 60
 in single-user mode, 374
 standard Samba during, 343
 step 1: POST process, 58–59
 step 2: BIOS-to-boot loader, 59–60
 step 3: kernel loading, 60–61
 step 4: `init` process, 61
 troubleshooting, 57–61, 373–374

boot.log file, 313

Bourne again shell. *See* bash

Bourne shell (`sh`), 84

`brctl` command, 326–328

bridge for wireless networking, 139, 325–329

broadband
 Internet connection setup, 146–148
 kinds of, 136, 138
 overview, 135

broken links, 105

browsers, 157–158. *See also* Firefox browser

burning CDs and DVDs
 data, with K3b, 236–238
 data, with Nautilus, 129–131
 with GParted program ISO, 30
 from ISO file, 130–131, 237–238
 Nautilus default action, 112, 125

• C •

`c` in `ls` listings, 111

cable modems, 136, 139. *See also* Internet connection

Calc (OpenOffice.org)
 documents in, 201
 formats supported, 201–202
 formatting bar, 200
 formula bar, 200–201
 GNOME icon for, 71
 menu bar, 197–198
 quick-launch icon for, 192
 standard toolbar, 199–200

capturing screenshots (GIMP), 249–250

case
 commands sensitive to, 66
 for environment variable names, 91
 file and directory names sensitive to, 100
 for variable names, 91

`cat` command, 185–186

`cd` command, 115–116

CD drive designations, 114

- CD player (GNOME), 228–230
 - CDDb Track Editor (GNOME), 229–230
 - CDs. *See also* burning CDs and DVDs
 - ejecting, 124
 - Fedora, damaged, 44
 - Fedora, getting from vendor, 42
 - Nautilus default action, 111–112
 - using in Nautilus, 124–126
 - Cedega (TransGaming), 259–261
 - Certified hardware, 34
 - character devices in `ls` listings, 111
 - `chmod` command, 354
 - `clear` command, 98
 - CLI (command line interface). *See also*
 - `bash` (Bourne again shell); commands
 - advantages of, 65, 83
 - autocompletion feature, 66, 89–90
 - booting to, 371–372
 - GUI desktops versus, 83
 - history feature, 66
 - in rescue mode, 63
 - switching to GUI environment, 370–371
 - terminal window for, 65–66
 - codecs, video, 241
 - CodeWeavers CrossOver Office,
 - 241, 258–259
 - colors
 - Firefox preferences for, 162
 - with `ls` command, 105–106
 - commands. *See also* `bash` (Bourne again shell); CLI (command line interface); *specific commands*
 - administration, 384
 - for archiving and compressing, 378–379
 - arguments with, 87
 - autocompletion feature, 89
 - `bash`, built-in, 379–380
 - as case-sensitive, 66
 - components of, 86
 - conventions in this book, 2–3
 - for file attributes, 381
 - file locators, 381–382
 - for file organization, 380–381
 - file viewers, 185–186, 382
 - filesystem, 115–116, 382–383
 - flags or options with, 86–87
 - for help, 378
 - `--help` option, 87
 - history feature, 66, 90
 - kernel module handling, 385
 - literal versus interpolated strings with, 94
 - `mttools`, 383–384
 - other terms for, 3
 - for processes, 385–386
 - redirecting output, 96–97
 - terminal window for, 65–66
 - wildcard expansion with, 88
 - Compatible hardware, 34
 - compressing files, 269–270, 378–379
 - connecting to the Internet. *See* Internet connection
 - connection types for printers, 289–290
 - cookies, 160–161
 - copying files and folders
 - `cp` command for, 116
 - in Nautilus, 121–122
 - copyleft software, 15
 - copyrights, 15–16
 - core dump files, 90
 - `cp` command, 116
 - CrossOver Office (CodeWeavers),
 - 241, 258–259
 - CUPS (Common Unix Printing System), 286
- D •
- `d` in `ls` listings, 108, 111
 - `dash` (-)
 - in file and directory names, 100
 - in permission sets, 108–110
 - Debian GNU/Linux, 16–17
 - deleting or removing
 - applets from GNOME panel, 75
 - cookies in Firefox, 160
 - directories in Nautilus, 122–123
 - directories with `rmdir`, 116
 - directory and files, 116
 - emptying trash can (GNOME), 76–77, 123
 - files in Nautilus, 122–123
 - Linux, 64
 - program from GNOME panel or desktop, 76
 - secondary panel (GNOME), 73
 - trash can for (GNOME), 72
 - unmounting floppy disks, 112
 - users, 284

- desktops. *See also* GNOME (GNU Network Object Model Environment)
 - booting to, 371–372
 - choosing when installing Fedora, 51
 - CLI versus, 83
 - GNOME as default, 67
 - GUI hung, 373
 - GUI won't start, 369
 - KDE, 51
 - switching environments in GNOME, 72
 - VT-7 as default for, 85
- /dev directory, 101, 114–115
- development of Linux, 12
- DHCP (Dynamic Host Configuration Protocol)
 - for broadband Internet connection, 148
 - for dialup Internet connection, 146
 - setup when installing Fedora, 48
 - for wireless Internet connection, 151
- dialup modems, 137, 142–146
- Digital Millennium Copyright Act (DMCA), 246
- Digital Subscriber Line (DSL), 136, 139. *See also* Internet connection
- directories. *See also* filesystem; permission sets; *specific directories*
 - for Amarok music player, 231
 - base, 100–102
 - changing with `cd`, 115–116
 - closing in Nautilus, 119
 - copying in Nautilus, 121–122
 - creating in Nautilus, 120–121
 - creating with `mkdir`, 116
 - deleting in Nautilus, 122–123
 - for file sharing, creating, 352–353
 - for log files, 311
 - moving in Nautilus, 121–122
 - naming conventions, 100
 - opening in Nautilus, 119
 - partitions versus, 113–115
 - removing with `rmdir`, 116
- DirectX, Cedega support for, 259
- disabling. *See* turning off
- Disk Druid (Red Hat), 47
- distributions of Linux, 16–18
- DivX video format, 241, 243
- DivX Web Player, 241
- DMCA (Digital Millennium Copyright Act), 246
- dollar sign (\$)
 - in `bash` prompt, 84–85
 - retrieving variable values, 93
- DOS, `mtools` commands for, 383–384
- dot (`.`)
 - in file and directory names, 100
 - starting hidden files, 86, 106
- double quotes (") for strings, 94
- downloading software, 276–277. *See also* Internet resources
- `Dr. Hardware` tool, 38
- Draw (OpenOffice.org)
 - described, 208
 - documents in, 211–212
 - drawing bar, 212
 - formats supported, 212
 - line and filling bar, 211
 - menu bar, 209–210
 - ruler, 211
 - standard toolbar, 210–211
- drive designations, 114
- drivers
 - in `ls` listings, 110–111
 - for `ndiswrapper`, 323
 - for 3-D video cards, 254
- DSL (Digital Subscriber Line), 136, 139. *See also* Internet connection
- `du` command, 86–87
- dual booting
 - assessing your situation, 20–21
 - defined, 19
 - installing Windows first, 21
 - Internet connection with, 136
 - options for, 19–20
 - partitioning from scratch for, 21
 - removing Linux from system, 64
 - resizing partitions for, 27–28
 - viewing Windows partitions, 22–27
 - Windows information needed for, 22
- DVD drive designations, 114
- DVD-ROM with this book
 - booting from, 389–390
 - contents, 387

damaged, getting replacement, 44
distributions covered on, 17
Fedora on, 39, 43, 387, 390
system requirements, 388–389
troubleshooting, 391
using, 389
using without a DVD drive, 388
DVDs. *See also* burning CDs and DVDs
ejecting, 124
Nautilus default action, 111–112
problems installing from disc you
burned, 44
using in Nautilus, 124–126
watching, 246
Dynamic Host Configuration Protocol.
See DHCP
dynamic IP address, 141

● E ●

echo command, 93–94
efficiency of Linux, 12
Ekiga Internet phone program, 179–181
emacs text editor, 186
e-mail. *See also* Evolution program
checking in Evolution, 173
creating and sending in Evolution,
171–173
enabling firewall for server, 54
emptying trash can (GNOME), 76–77, 123
enabling
Java and JavaScript in Firefox, 162
servers, firewall settings for, 54
3-D video card drivers, 255
encryption, Firefox preferences for, 163
environment variables
commonly used (table), 91–92
creating and exporting, 95
in interpolated strings, 94
as part of shell environment, 87
uppercase convention for, 91
variables versus, 91
equals sign (=) for assigning variable
values, 92
equation editor. *See* Math (OpenOffice.org)

error messages
[FAILED] status for services, 61
for hardware failure when booting, 58
for Internet connection errors, 151
during kernel loading, 60–61
kernel panic, 60–61
log file for, 312–313
operating system not found, 60
printer policy for, 293
searching Google for, 61
/etc directory, 101, 102
/etc/group file, 109
/etc/passwd file, 109
/etc/profile file, 95
Ethernet cards. *See* NICs (network
interface cards)
Evanescence site, 230
Evolution program
creating and sending e-mail, 171–173
described, 165
GNOME icon for, 71
setting up, 166–171
Evolution Setup Assistant
Account Management screen, 170–171
checking for e-mail, 173
Identity screen, 166–167
Receiving E-Mail screen, 167–168
Receiving Options screen, 168–169
Sending E-Mail screen, 169–170
Time Zone screen, 171
Welcome screen, 166
execute (x) permission, 109–110
exit command, 85
export command, 95
ext2 filesystem, 113, 128
ext3 filesystem, 113, 128
extensions (Firefox), 163–164, 242–243
external modems, 139

● F ●

[FAILED] status for services, 61
FAQ for Fedora, 39, 246, 274, 276
fdisk command, 112, 127
FDISK tool (Windows 98), 25–27
fdisk/mbr command, 64

Fedora. *See also* installing Fedora; *specific components*

Audio Configuration tool, 227–228

on the DVD-ROM, 39, 43, 387, 390

first boot after installing, 53–56

NIC detection by, 147

rebooting, 63

reference sites, 39

as RHEL proving ground, 18

shutting down, 63–64

as this book's focus, 17–18, 42

unique groups for users in, 110

updating, 270–271

upgrading, 277

Fedora Documentation Project site, 39

Fedora Forum, 39, 276

Fedora News site, 39

Fedora Package Manager

adding repositories, 274–275

adding the software in this book, 275–276

finding software, 272–274

opening, 272

Fedora Printer Configuration tool

adding a printer, 288–291, 337–339

modifying printer properties, 291–295

opening, 286

overview, 286–287

print server settings, 288, 339–341

Fedora Project

sending hardware profile to, 55–56

Web site, 17, 41

Wiki site, 39

file formats

Calc support for, 201–202

DOS, mtools commands for, 383–384

for downloaded files, 268–269

Draw support for, 212

Impress support for, 208

for video, 243–244

video players' support for, 240–241

Windows, finding programs for, 257

Writer support for, 196–197

file manager. *See* Nautilus file manager (GNOME)

file server. *See* Samba file server

File Transfer Protocol (FTP),

54, 178–179, 306

filenames

autocompletion feature, 66, 89–90

hidden files, 86, 106

Linux conventions for, 100

wildcard expansion for, 88

files. *See also* directories; file formats;

filenames; `ls` command

closing in Nautilus, 119

commands for archiving, 378–379

commands for attributes, 381

commands for locating, 381–382

commands for organization, 380–381

commands for viewing, 185–186, 382

compressing, 269–270, 378–379

copying in Nautilus, 121–122

copying with `cp`, 116

creating in Nautilus, 120–121

deleting in Nautilus, 122–123

deleting with `rm`, 116

DOS, mtools commands for, 383–384

downloaded, opening, 267–269

hidden, 86, 106

information in `ls` long format, 107

links to, 105

moving in Nautilus, 121–122

moving with `mv`, 116

opening in Nautilus, 119, 121

packaging and compressing, 269–270

redirecting command output to, 96

renaming with `mv`, 116

SUID (SetUID) permission for, 106

text, 185–190

types, 110–111

file-sharing programs, 184. *See also* Samba

file server

filesystem. *See also* directories; `ls`

command

base directory standards, 100

commands for, 115–116, 382–383

defined, 47

monitoring, 299–300

naming conventions, 100

navigating in Fedora, 111–113

root directory, 47, 100–102

types, 113, 128

usage in this book, 100

Windows compared to Linux, 99–100

- `find` command, 90
- finding. *See* searching
- Firefox browser
 - adding plug-ins, 241–242
 - Advanced preferences, 163
 - based on Mozilla, 158
 - changing font size, 162
 - checking plug-ins installed, 240–241
 - Content preferences, 161–162
 - on the DVD-ROM, 387
 - extensions, 163–164, 242–243
 - General preferences, 158–159
 - GNOME icon for, 71
 - home page, 158–159
 - opening Preferences window, 158
 - plug-ins, 163, 240–242
 - Privacy preferences, 159–161
 - RSS extensions, 184
 - surfing the Web with, 164–165
 - tabbed browsing feature, 164–165
 - Web site, 158
- firewall
 - configuring for networking, 305–307
 - defined, 53
 - need for, 53–54
 - print server settings, 307, 344
 - setup on first boot, 53–55
- `fixmbr` command (XP), 64
- flags for commands, 3, 86–87
- Flash (Macromedia), 230, 275
- floppy drives
 - designations for, 114
 - formatting disks, 129
 - using in Nautilus, 112, 125
- folders. *See* directories; filesystem
- fonts, Firefox settings for, 162
- forgotten root password, 374
- formula editor. *See* Math (OpenOffice.org)
- forum for Fedora, 39, 276
- forward slash (/) for directories, 100
- free software
 - described, 9
 - GPL protecting, 10, 14–15
 - Linux as, 12
- FSF (Free Software Foundation), 14
- FTP (File Transfer Protocol),
 - 54, 178–179, 306
- **G** •
 - games, 3-D. *See* 3-D games
 - games on the DVD, 387
 - `gedit` text editor, 188–190
 - General Public License (GPL), 10, 14–15, 53
 - gFTP program, 178–179
 - GIMP graphics program
 - basic tools, 248–249
 - capturing screenshots, 249–250
 - editing image files, 250–252
 - effects, 252–253
 - opening images, 250
 - opening the program, 247
 - scripts, 252–253
 - setup routine, 247
 - sources for images, 248
 - toolboxes, 248–249
 - undoing effects, 253
 - GNOME (GNU Network Object Model Environment). *See also specific components*
 - adding program to panel or desktop, 75–76
 - booting to, 371–372
 - CD player, 228–230
 - choosing when installing Fedora, 51
 - as default desktop, 67
 - desktop icons, 76–77
 - on the DVD-ROM, 387
 - further information, 67
 - GUI won't start, 369
 - hiding/restoring applications, 72
 - Lock Screen option, 70, 85
 - menus, 67–69
 - moving programs between workspaces, 72
 - panels, customizing, 73–76
 - panels, overview, 71–72
 - removing program from panel or desktop, 76
 - switching desktop environments, 72
 - switching to CLI, 370–371
 - taskbar, 72
 - themes, 77–79
 - GNOME Office suite, 192
 - GNU (GNU's not Unix), 14
 - Google, 34, 61
 - Gookin, Dan (*PCs For Dummies*), 387

GParted program
 altering partitions with, 30–33
 backing up before using, 30
 burning the ISO onto CD, 30
 downloading, 29–30
 GPL (GNU General Public License),
 10, 14–15, 53
 graphics
 capturing screenshots, 248–249
 Draw (OpenOffice.org) for, 208–212
 editing image files, 250–252
 GIMP for, 247–253
 greater than sign (>) as redirection
 operator, 96
 groups
 adding, 284–285
 for file server access, 351–352, 358–359
 in `ls` listings, 107
 modifying, 285–286
 overview, 110
 permissions for, 109–110
 GRUB (GRand Unified Boot loader),
 47, 59–60

• **H** •

hacker, defined, 13
Hackers (Levy), 13
 halt command, 63
 hard drives
 adding second for Linux, 19, 28
 designations for, 114–115
 determining free space on, 27–28
 monitoring, 299–300
 options when installing Fedora, 46–47
 resizing partitions for dual booting, 27–28
 space needed for Linux, 22–24, 27, 29
 space needed for Windows, 28
 viewing with `fdisk`, 112, 127
 hard links, 105
 hardware. *See also specific kinds*
 Certified, Supported, or Compatible, 34
 checking compatibility with Linux, 34–39
 hardware-detect tools, 38
 Internet connection setup, 139
 laptop compatibility information, 35
 POST problems with, 58–59
 Win prefix in names, 38–39
 Helix DNA program, 240

help
 commands for, 378
 --help option for commands, 87
 man pages for, 97, 378, 380
 hidden files, 86, 106
 hiding/restoring programs in GNOME, 72
 history command, 90
 history feature, 66, 90, 92
 HISTSIZE environment variable, 91–92
 home directory, 86
 HOME environment variable, 91
 Home icon
 GNOME panel, 76
 Nautilus file manager, 117–118
 host name setup, 48
 HTTP and HTTPS services, firewall
 settings for, 306

• **I** •

icons
 GNOME desktop, 72, 76–77
 in margins of this book, 6
 on panels (GNOME), 71
 quick-launch, for OpenOffice.org, 192
 IDE drive designations, 114–115
 IM (instant messaging), 173–176
 images. *See* graphics
 Impress (OpenOffice.org)
 documents in, 207
 drawing bar, 207–208
 formats supported, 208
 GNOME icon for, 71
 line and filling bar, 206–207
 menu bar, 204–205
 Presentation Wizard, 202–204
 quick-launch icon for, 192
 standard toolbar, 205–206
 view switchers, 207
 init program, 61
 installing. *See also* installing Fedora
 Amarok music player, 230–231, 275
 BitTorrent, 177
 bridging software, 326
 Firefox extensions, 242–243
 Firefox plug-ins, 242
 Java, 275
 kmod-ndiswrapper package, 322–323
 K3b program, 129, 275

- Mplayer, 275
- OpenOffice.org suite, 192
- RealPlayer, 275–276
- Samba, 341, 355
- Skype, 182–183
- software using Fedora Package Manager, 272–276
- themes (GNOME), 78–79
- 3-D video card drivers, 254
- Thunderbird, 275
- Windows, replacing Linux, 64
- Windows before Linux, 21
- xine video player, 242
- installing Fedora. *See also* preparing to install Linux
 - backing out before finishing, 42
 - booting first time after, 53–56
 - checking distribution burns, 367–368
 - damaged CDs or DVDs, 44
 - differences in screens for, 42
 - downloading CD images for, 41–42
 - DVD-ROM with this book for, 41, 43, 389–390
 - frozen installer, 366–367
 - getting CDs from vendor, 42
 - GUI won't start, 369
 - information erased during, 20
 - with no mouse, 43–44
 - options for, 19–20
 - reinstalling for problems with booting, 60
 - steps for, 43–53
 - text mode for, 42–44
 - troubleshooting, 366–369
 - video card fails, 368–369
- instant messaging (IM), 173–176
- Integrated Services Digital Network (ISDN), 138–139. *See also* Internet connection
- internal modems, 139
- Internet connection
 - broadband options, 135–136, 138
 - configuring for broadband, 146–148
 - configuring for dialup, 141–146
 - configuring for wireless, 148–151
 - dialup, 137
 - firewall for, 53–54
 - hardware setup, 139
 - information needed from ISP, 140–141
 - ISDN, 138
 - selecting an ISP, 140
 - troubleshooting, 151–153, 176
 - useful sites, 153–154
 - wireless, 138
- Internet phone
 - Ekiga for, 179–181
 - KPhone for, 179
 - Linphone for, 179
 - overview, 179
 - Skype for, 179, 181–183
- Internet radio, 233–234
- Internet Relay Chat (IRC) programs, 184
- Internet resources
 - Amarok music player, 230–231
 - base directory standards, 100
 - basic recommendations, 153–154
 - BitTorrent, 176–178
 - Cedega, 259
 - CodeWeavers CrossOver Office, 258
 - distributions of Linux, 17–18
 - Dr. Hardware tool, 38
 - Evanesence site, 230
 - Fedora CD images, 41–42
 - Fedora FAQ, 39, 246, 274, 276
 - Fedora Forum, 39, 276
 - Fedora installation package, 41–42
 - Fedora reference sites, 39
 - Firefox extensions, 242
 - Firefox plug-ins, 241
 - Firefox site, 158
 - GNU history, 15
 - GParted download, 29
 - hardware compatibility information, 34–35
 - Internet phone programs, 179
 - Konqueror browser, 158
 - LinModem project, 137
 - MathML format information, 224
 - music downloads, 230–233
 - ndiswrapper site, 148, 322
 - Nero Burning ROM program, 30
 - nonapproved Fedora packages, 254
 - PuTTY SSH client, 308
 - RHEL hardware compatibility list, 34
 - security reference sites, 314
 - Skype client, 179
 - software downloads, 276–277
 - themes (GNOME), 77
 - 3-D game repositories, 255
 - 3-D video card drivers, 254
 - time server, 55

Internet resources (*continued*)

- troubleshooting help, 151
- Vista support for Partition Magic, 20
- VMware products, 261
- Wikipedia, 232
- Wiley Product Technical Support, 391
- Windows/Linux programs, 257
- Wine, 262–263
- Win4Lin, 259
- wireless networking information, 138–139, 148
- Internet Service Provider (ISP), 139–141.
 - See also* Internet connection
- interoperability of Linux, 11
- interpolated strings, 94
- IP address, 141, 148
- IP telephony. *See* Internet phone
- ip6tables service, turning off, 304
- iptables service, 304
- IRC (Internet Relay Chat) programs, 184
- ISDN (Integrated Services Digital Network), 138–139. *See also* Internet connection
- ISO files
 - burning with K3b, 130–131, 237–238
 - for GParted program, 30
- ISP (Internet Service Provider), 139–141.
 - See also* Internet connection

• **J** •

- Java, 162, 275
- JavaScript, enabling in Firefox, 162
- job options for printers, 294–295

• **K** •

- KDE (K Desktop Environment), 51, 387
- kernel
 - defined, 15
 - differences among distributions, 18
 - failure loading, 60–61
 - loading during boot, 60
 - module-handling commands, 385
 - recovering from update problems, 61
 - Torvalds as benevolent dictator of, 15
- kernel panic message, 60–61
- keyboard, installation options for, 45–46

- KIRC program, 184
- Knoppix distribution, 16, 17
- KOffice suite, 192
- Konqueror browser, 158
- KPhone Internet phone program, 179
- K3b program
 - burning ISO file to CD or DVD, 130–131, 236–238
 - closing, 238
 - installing, 129, 275
 - opening, 236
- kudzu service, turning off, 304

• **L** •

- 1 in `ls` listings, 111
- language options during installation, 44–45
- laptop compatibility, checking, 35
- LeBlanc, Dee-Ann (*Linux For Dummies*), 1–6
- `less` command, 186
- Levy, Steven (*Hackers*), 13
- LILO (Linux LOader), 59–60
- LimeWire file-sharing program, 184
- links, 105, 111
- LinModem project, 137
- Linphone Internet phone program, 179
- Linspire distribution, 16–17
- Linux All-in-One Desk Reference For Dummies* (Barkakati), 1
- Linux and Windows on same machine.
 - See* dual booting; Windows programs under Linux
- Linux For Dummies* (LeBlanc and Blum).
 - See also* DVD-ROM with this book
 - assumptions about the reader, 3–4
 - Fedora as focus of, 17–18
 - icons in margins, 6
 - organization, 4–5
 - overview, 1–2
 - using, 2–3, 6
- LinuxQuestions.org site, 39
- literal text strings, 94
- live distributions, 17, 20
- Lock Screen option, 70, 85
- log files, 311–313
- logging in to `bash` shell, 85
- logging out from VTs, 85
- lowercase. *See* case

`ls` command
 colors displayed with, 105–106
 file size in listings, 107
 file types in listings, 110–111
 groups in listing, 107
 listing all home directory files, 86
 listing hidden files, 106
 listing only `.doc` files, 88
 long format listing, 107
 overview, 104–106
 owners in listing, 107
 permission sets in listings, 107–109
 redirecting output, 96–97

• M •

Macintosh OS X, SSH with, 311
 Macromedia Flash, 230, 275
`mahjongg` command, 88
 mail server, firewall settings for, 54, 306
`MAILCHECK` environment variable, 92
`maillog` file, 313
`man` pages, 97, 378, 380
 Mandriva distribution, 17
 mapping
 to shared folders, 362
 Windows workstation to Samba print server, 346
 Math (OpenOffice.org)
 described, 218–219
 document section, 223
 formats supported, 223–224
 menu bar, 219–220
 standard toolbar, 220–221
 MathML (MML) format, 223–224
 MBR (Master Boot Record), 59
`mdmonitor` service, turning off, 304
`/media` directory, 103
 menus. *See also* menus (GNOME)
 Base (OpenOffice.org), 215–216
 Calc (OpenOffice.org), 197–198
 Draw (OpenOffice.org), 209–210
 GIMP, 248
 Impress (OpenOffice.org), 204–205
 Math (OpenOffice.org), 219–220
 Totem, 245–246
 User Manager tool, 280
 Writer (OpenOffice.org), 193–194

menus (GNOME)
 adding buttons to panel, 75–76
 Applications, 67–68
 Places, 67–69
 shortcut or context menus, 73, 74
 submenus, 68
 System, 68–69
 messages log file, 312–313
 Microsoft Office
 CrossOver Office for, 241, 258–259
 formats supported by OpenOffice.org,
 196–197, 201–202, 208
 Minix, 12
`mkdir` command, 116
 MML (MathML) format, 223–224
`/mnt` directory, 101
 modems. *See also* Internet connection
 cable, 136, 139
 dialup, 137, 142–146
 DSL, 136, 139
 external versus internal, 139
 ISDN, 139
 LinModem project, 137
 satellite, 136, 139
 WinModems, 38–39, 137
 monitoring the system. *See* System Monitor
`more` command, 186
 mount point, 47, 113
 mouse, installing Fedora without, 43–44
 MOV video format, 241, 244
 moving files and folders
 with `mv`, 116
 in Nautilus, 121–122
 Mozilla program, 158, 241, 242
 MPG (Moving Pictures Experts Group)
 format, 240, 244
 Mplayer plug-in, 233, 275
`mtools` commands, 383–384
 multiplatform, defined, 11
 multiprocess, defined, 11
 multiuser, defined, 11
 music. *See* music CDs; sound
 music CDs. *See also* CDs
 adding to database, 229–230
 Nautilus default action, 112, 124, 126
 ripping tracks from, 234–236
 switching to GNOME CD player, 228–229
 Totem player, 228
`mv` command, 116

• **N** •

names and naming. *See also* filenames;
passwords

drive designations, 114–115

filesystem conventions, 100

`mv` command for renaming files, 116

variables versus environment
variables, 91

`nano` text editor, 186–187

Nautilus file manager (GNOME)

accessing local Windows drives, 127–128

accessing network drives, 128

burning CDs or DVDs in, 125

changing permissions, 123–124

closing folders, 119

closing windows, 119

copying files and folders, 121–122

creating files and folders, 120–121

default actions, 111–112

deleting files and folders, 122–123

emptying trash can, 123

file browser, 118

formatting disks, 129

Home icon, 117–118

mounting partitions, 112–113, 127

moving files and folders, 121–122

opening files, 119, 121

opening folders, 119

packaging and compressing files, 269–270

preferences, 125–126

running programs, 121

Select Places window, 119

Tree view, 119

using removable media in, 124–127

View menu commands, 119–120

`ndiswrapper`, 148, 322–323

Nero Burning ROM program, 30

Network Time Protocol (NTP) setup, 55

networking. *See also* NICs (network
interface cards); Samba file server;
Samba print server; wireless
networking

accessing drives in Nautilus, 128

controlling services, 303–305

firewall settings, 53–55, 305–307

monitoring, 299

setup when installing Fedora, 47–49

turning off unneeded services, 304–305

news site for Fedora, 39

NFS4 (Network File Sharing), 306

NICs (network interface cards)

configuring wireless as WAP, 324–325

detected automatically by Fedora, 147

device names for, 147

installing wireless, 321–322

Internet connection setup, 147–148

`ndiswrapper` for wireless, 148, 322–323

required for broadband connections, 146

wireless card setup, 149–151, 320–321

`nmbd` service, 343

Novell SuSE distributions, 17–18, 110

`ntfs` filesystem, 113, 128

NTP (Network Time Protocol) setup, 55

NVIDIA binary drivers, 254

• **O** •

Office (Microsoft)

CrossOver Office for, 241, 258–259

formats supported by OpenOffice.org,
196–197, 201–202, 208

office suites, 191–192. *See also*

OpenOffice.org suite

open source approach, 13–16

OpenDocument formats

Calc support for, 201

Draw support for, 212

Impress support for, 208

Math support for, 223

Writer support for, 196

OpenOffice.org suite. *See also specific
packages*

accessing packages, 192

Base package, 212–218

Calc package, 71, 197–202

default packages with Fedora, 192

Draw package, 208–212

on the DVD-ROM, 387

GNOME icons for, 71

Impress package, 71, 202–208

installing, 192

Math package, 218–224

quick-launch icons, 192

Writer package, 71, 193–197

operating system not found
 message, 60
 operating systems, 10–12
 options for commands, 3, 86–87
 owners, 107, 109–110. *See also* users

• p •

package managers
 defined, 18
 differences among distributions, 18
 using Fedora's, 272–276
 Package Updater, 270–271
 packaging files, 269–270
 panels (GNOME)
 customizing, 73–76
 overview, 71–72
 parameters for commands, 3, 86–87
 Partition Magic program, 20, 29
 partitions and partitioning
 backing up first, 29–30
 commercial programs for, 20, 29
 determining hard drive free space, 27–28
 directories versus partitions, 113–115
 GParted program for, 29–34
 monitoring, 299–300
 mount point, 47, 113
 mounting partitions in Nautilus,
 112–113, 127
 noting for Windows, 21
 options when installing Fedora, 46–47
 resizing partitions for dual booting,
 27–28
 from scratch for dual boot, 21
 space needed for Linux, 22–24, 27, 29
 space needed for Windows, 29
 viewing in Vista, 22–23
 viewing in Windows 98, 25–27
 viewing in Windows 2000 and XP, 23–24
 viewing with `fdisk`, 112
 Windows information needed for, 22
 passwords
 for dialup Internet connection, 145
 forgotten root, 374
 with Lock Screen option, 70
 never giving out, 302
 root, 50, 112, 374
 secure, choosing, 302
 for VTs, 85
 PATH environment variable, 92
PCs For Dummies (Gookin), 387
 pcsd service, turning off, 305
 permission sets
 changing in Nautilus, 123–124
`chmod` command for, 354
 defined, 108
 for file sharing, 352–354
 in `ls` listings, 107–108
 SGID (Set Group ID), 106
 SUID (SetUID), 106
 triplets in, 108–109
 pictures. *See* graphics
 Pidgin IM program, 174–176
 ping command, 152–153
 pipe symbol (`|`) for redirecting output, 97
 platforms supported by Linux, 11
 playing
 downloaded music, 230–233
 DVDs, 246
 Internet radio, 233–234
 3-D games, 253–256
 Web clips, 240–241, 243
 plug-ins
 adding to Firefox, 241–242
 checking in Firefox, 240–241
 gedit text editor, 188–189
 security problems from, 163
 policies for printers, 292–293
 popups, blocking in Firefox, 162
 portability of Linux, 12
 ported, defined, 12
 POST (Power-On Self-TEST), 58–59
 power problems, 64
 PPD (Postscript Printer Description) file,
 290–291, 293
 preemptive multitasking, 11
 preparing to install Linux. *See also* dual
 booting
 adding second hard drive, 19, 28
 bookmarking Fedora sites, 39
 checking hardware compatibility, 34–39
 working with partitions, 29–34
 presentation program. *See* Impress
 (OpenOffice.org)

print server. *See* Samba print server
printers

- adding, 288–291
- connection types, 289–290
- controlling state of, 292
- CUPS for, 286
- job options, 294–295
- modifying properties, 291–295
- policies for, 292–293
- PPD files for, 290–291, 293
- restricting access to, 293

processes, 296–298, 385–386

programs. *See also* commands; *specific programs*

- adding to GNOME panel or desktop, 75
- boot-loaders, 59–60
- choosing when installing Fedora, 51–52
- for downloaded files, 268–269
- DVD burners, 30
- executable permission for, 111
- Fedora Package Manager for installing, 272–276

file sharing, 184

finding, 272–273, 276–277

on GNOME taskbar, 72

hardware-detect tools, 38

hiding/restoring in GNOME, 72

Internet phone, 179

IRC, 184

moving between workspaces in GNOME, 72

mttools suite, 383–384

other terms for, 3

for partitioning, 20, 29

recognizing in Linux, 111

RSS readers, 184

running from Nautilus, 121

for running Windows software, 258–264

as security holes, 311

starting from the shell, 87–88

updating securely, 302–303

Windows-equivalents, 257

prompt

for bash shell, 84–85

PS1 environment variable for, 92–93

resetting in bash, 98

proxy servers, Firefox setup for, 163

PS1 environment variable, 92–94

PuTTY SSH client, 308

• Q •

QT video format, 241, 244

question mark (?) as wildcard, 88

QuickTime player (Apple), 241

• R •

r (read) permission, 109

RealPlayer, 233, 275–276

reboot command, 63

Red Hat Enterprise Linux (RHEL), 17–18, 34

Red Hat Package Manager (RPM) on the DVD, 387

redirecting output, 96–97

reiserfs filesystem, 113, 128

Remember icon, 6

removable media. *See also specific media*
defined, 124

formatting in Nautilus, 129

Nautilus preferences for, 125–126

using in Nautilus, 124–127

rescue mode, 62–63

reset command, 98

resizing partitions for dual booting, 27–28

resolution, changing for screen, 372–373

restoring/hiding programs in GNOME, 72

restricting access to printers, 293

RHEL (Red Hat Enterprise Linux), 17–18, 34

ripping music tracks from CDs, 234–236

rm command, 116

RM video format, 244

rmdir command, 116

root directory

base directories in, 101–102

defined, 47, 100

root user versus, 100

root password setup, 50

root user

file ownerships for, 109–110

Lock Screen option not available for, 70

root directory versus, 100

su command for, 112, 127

for system administration, 280

router for wireless networking, 139

RPM (Red Hat Package Manager) on the DVD, 387

RSS readers, 184

• S •

Samba file server

- configuring Samba, 356–358
- creating folders to share, 352–353
- creating Linux groups, 351–352
- creating Linux user accounts, 358–359
- creating Samba user accounts, 359–361
- creating shared folders in Samba, 361–362
- installing Samba, 355
- mapping to shared folders, 362
- planning a sharing policy, 350–351
- Samba on the DVD-ROM, 387
- Samba overview, 350
- setting folder permissions, 353–354
- setting group permissions, 352–353
- starting Samba, 355–356
- troubleshooting, 362–363

Samba print server

- allowing guests to print, 345
- configuring Samba, 341–342
- connecting local printer, 337
- firewall settings, 307, 344
- installing Samba, 341
- mapping Windows workstation to, 346
- planning printer sharing, 336–337
- Printer Configuration settings, 288, 337–341
- Samba defined, 336
- Samba on the DVD-ROM, 387
- starting Samba, 342–343
- troubleshooting, 346–347

satellite modems, 136, 139. *See also*

Internet connection

- scalability of Linux, 11
- screen resolution, changing, 372–373
- screensaver, 70, 85
- screenshots, capturing with GIMP, 249–250

scripts. *See also* programs

- GIMP, 252–253
- other terms for, 3
- shell scripting defined, 95
- as text files, 185

SCSI drive designations, 114–115

searching

- computer with Beagle, 131–132
- for error message information, 61, 151
- file locator commands, 381–382
- for hardware compatibility information, 34
- for software, 272–273, 276–277

Second Life 3-D game, 256

secure log file, 313

Secure Shell. *See* SSH

security

- importance of, 301
- log files for, 311–313
- networking issues, 303–307
- password selection for, 302
- plug-in issues, 163
- SELinux setup on first boot, 55
- software holes, 311
- SSH for, 307–311
- updating software securely, 302–303
- Web sites, 314
- for wireless networking, 330–331

SELinux (Security-Enhanced Linux) setup, 55

sendmail service, turning off, 305

services

- boot problems from, 61
- defined, 61
- Fedora tool for, 303–304
- network, firewall settings for, 306–307
- starting Samba, 342–343
- turning off unneeded, 304–305

Session Internet Protocol (SIP), 179

SetUID (SUID) permission, 106

SGID (Set Group ID) permission, 106

sh (Bourne shell), 84

shell scripting. *See* scripts

shells. *See also* bash (Bourne again shell)

- built-in commands, 379–380
- defined, 84
- environment variables with, 87
- help commands, 378
- starting a session, 84–85
- starting programs from, 87–88
- VTs for, 85

shutting down Linux, 63–64

single quotes (') for strings, 94

single-user mode, 374

SIP (Session Internet Protocol), 179

Skype Internet phone program, 179, 181–183

Slackware distribution, 17

smartd service, turning off, 305

smbd service, 343

soft links, 105

software. *See* programs

- sound. *See also* music CDs
 - Internet radio, 233–234
 - playing downloaded music, 230–233
 - troubleshooting, 225–228
 - volume control, 71, 226–227
- sound card, 56, 225–228
- Sound Juicer CD Extractor, 234–236
- spamassassin service, turning on, 305
- spreadsheet program. *See* Calc (OpenOffice.org)
- SSH (Secure Shell)
 - firewall settings for, 54, 307
 - installed by default, 307
 - Linux-to-Linux connections, 310
 - Mac OS X-to-Linux connections, 311
 - overview, 307
 - Windows client software for, 307–308
 - Windows-to-Linux connections, 308–310
- stability of Linux, 12
- Stallman, Richard (GNU founder), 13–15
- StarOffice suite, 192
- starting Linux. *See* booting
- static IP address, 141
- static networking setup, 48–49
- `stderr` (standard error), 95
- `stdin` (standard input), 95
- `stdout` (standard output), 95
- stopping Linux, 63–64
- Straw RSS reader, 184
- streaming audio (Internet radio), 233–234
- strings, literal versus interpolated, 94
- `su` command, 112, 127
- SUID (SetUID) permission, 106
- Supported hardware, 34
- SuSE distributions (Novell), 17–18, 110
- swap partition setup, 47
- switches for commands, 3, 86–87
- switching
 - between CLI and GUI environment, 370–371
 - environments in GNOME, 72
 - to text environment, 85
 - from Totem to GNOME CD player, 228–229
 - users, GNOME applet for, 71
 - views in Impress, 207
 - between VTs, 85

- system administration
 - commands for, 384
 - monitoring the system, 295–300
 - need for, 279
 - for printing, 286–295
 - for users and groups, 280–286
- System Log Viewer, 312–313
- System Monitor
 - File Systems tab, 296, 299–300
 - opening, 295
 - Process table columns, 297–298
 - Processes tab, 296–298
 - Properties options, 298
 - Resources tab, 296, 298–299
 - System tab, 296
- system requirements for the DVD, 388–389

● T ●

- tarballs, creating in Nautilus, 269–270
- Technical Stuff icon, 6
- telnet program, turning off, 307
- terminal window. *See also* VTs (virtual terminals)
 - for command-driven tasks, 65–66
 - starting a shell session, 84
- text editors
 - emacs, 186
 - gedit, 188–190
 - nano, 186–187
 - vi, 186
 - word processors versus, 188
- text environment, switching to, 85
- text files
 - commands for viewing, 185–186, 382
 - editing with gedit, 188–190
 - editing with nano, 186–187
 - joining with cat, 185–186
 - Linux uses for, 185
- 3-D games
 - Cedega for, 259–261
 - Linux challenges for, 253–254
 - Linux video card drivers for, 254–255
 - online, 255–256
 - repositories, 255
 - Second Life, 256
- Thunderbird, installing, 275

- tilde (~) for home directory, 86
 - time server setup, 55
 - Tip icon, 6
 - toolbars
 - Base (OpenOffice.org), 216
 - Calc (OpenOffice.org), 199–201
 - Draw (OpenOffice.org), 210–212
 - Impress (OpenOffice.org), 205–208
 - Writer (OpenOffice.org), 194–196
 - tools. *See* commands; programs
 - Torvalds, Linus (Linux creator), 1, 10, 12, 15
 - Totem movie/music player
 - described, 240
 - full version, 244, 246
 - menus, 245–246
 - starting, 244
 - switching to GNOME CD player, 228–229
 - video formats supported, 240–241
 - watching DVDs with, 246
 - TransGaming Cedega, 259–261
 - trash can (GNOME)
 - desktop icon for, 76–77
 - emptying, 76–77, 123
 - panel shortcut for, 72
 - removing desktop icon using, 76
 - troubleshooting
 - basic techniques, 365–366
 - BIOS problems, 59
 - boot problems, 57–61, 373–374
 - booting into wrong environment, 370–372
 - boot-loader problems, 60
 - checking distribution burns, 367–368
 - confusion in first-time users, 369–370
 - damaged installation CDs or DVDs, 44
 - DVD-ROM with this book, 390
 - frozen Linux installer, 366–367
 - GUI hung, 373
 - GUI won't start, 369
 - hang during boot, 373–374
 - IM connections, 176
 - init program, 61
 - Internet connection, 151–153
 - kernel update problems, 61
 - kernel-loading failure, 60–61
 - Linux installation, 366–369
 - POST problems, 58–59
 - rebooting after power problems, 64
 - rescue mode for, 62–63
 - root password forgotten, 374
 - Samba file server, 362–363
 - Samba print server, 346–347
 - screen resolution, changing, 372–373
 - service [FAILED] status, 61
 - sound card, 225–228
 - video card, 368–369
 - wireless networking, 333–334
 - turning off
 - computer, 63–64
 - Lock Screen option, 70
 - services not needed, 304–305
 - telnet, 307
 - turning on. *See* enabling
 - Tux, 10
- u •
- Ubuntu distribution, 16, 17
 - underscore (_) in file and directory
 - names, 100
 - uninstalling Linux, 64
 - UNIX, development of, 12
 - unmounting floppy disks, 112
 - Unofficial Fedora FAQ, 39, 246, 274, 276
 - updating
 - Fedora, 270–271
 - Firefox preferences for, 163
 - GNOME applet for, 71
 - kernel, recovering from problems after, 61
 - security issues for, 302–303
 - upgrading Fedora, 277
 - uppercase. *See* case
 - UPS (Uninterruptible Power Supply), 64
 - URL conventions in this book, 2–3
 - USB storage devices, 111
 - User Manager tool
 - adding new groups, 284–285
 - adding new users, 281–283
 - controlling behavior of, 281
 - deleting users, 284
 - modifying groups, 285–286
 - modifying users, 283–284
 - opening, 280
 - overview, 280

users. *See also* groups
 adding, 281–283, 358–359
 creating accounts for file sharing, 358–359
 deleting, 284
 modifying accounts, 283–284
 owners, 107, 109–110
 Samba accounts for, 358–359
 switching, GNOME applet for, 71
 User Manager tool for, 280–284
 users group, 110
 /usr directory, 101, 103–104
 utilities. *See* commands; programs

• U •

variables
 assigning values to, 92
 in bash shell, defined, 90
 case conventions for, 91
 environment variables versus, 91
 retrieving values from, 93
 vfat filesystem, 113, 128
 vi text editor, 186
 video. *See also* graphics; video CDs
 and DVDs
 adding plug-ins to Firefox, 241–242
 Apple formats, 241
 checking plug-ins installed in Firefox,
 240–241
 codecs, 241
 file formats, 243–244
 Firefox extensions for, 242–243
 Microsoft formats, 241
 players for, 240–241
 playing 3-D games, 253–256
 viewing movie files, 243–246
 Web clips, 239–243
 video card
 getting information at startup, 38
 3-D, Linux drivers for, 254
 troubleshooting, 368–369
 video CDs and DVDs. *See also* CDs; DVDs
 Nautilus default action, 112, 124, 126
 watching DVDs, 246
 video resolution, changing, 372–373
 virtual Linux machine, 20

virtual Windows machine, 20. *See also*
 Windows programs under Linux
 Vista. *See* Windows Vista
 VMware, 20, 261–262
 VMware Player, 261
 VMware Server, 261
 VoIP. *See* Internet phone
 volume control
 Amarok music player, 233
 GNOME icon for, 71
 using (GNOME), 226–227
 VTs (virtual terminals). *See also* terminal
 window
 clearing the screen, 98
 default for GUI, 85
 exiting, 85
 not locked by Lock Screen tool, 85
 resetting the prompt, 98
 starting a shell session, 84–85
 switching between, 85

• W •

w (write) permission, 109
 WAP (wireless access point). *See also*
 wireless networking
 configuring wireless NIC as, 324–325
 defined, 317
 as master device, 318
 purchasing, 317
 troubleshooting, 333–334
 Warning! icon, 6
 Web browsers, 157–158. *See also* Firefox
 browser
 Web clips
 adding Firefox extensions for, 242–243
 adding plug-ins for, 241–242
 checking plug-ins for, 240–241
 xine video player for, 242
 Web server, enabling firewall for, 54
 Web sites. *See* Internet resources
 wiki for Fedora Project, 39
 Wikipedia site, 232
 wildcards with commands, 88
 Wiley Product Technical Support, 391
 windows, closing in Nautilus, 119

- Windows and Linux on same machine. *See*
 - dual booting; Windows programs under Linux
 - Windows client software for SSH, 307–308
 - Windows Media Video (WMV) format, 240–241, 244
 - Windows 98
 - determining hard drive free space, 28
 - finding hardware information, 36
 - viewing partitions, 25–27
 - Windows programs under Linux
 - Cedega program for, 259–261
 - CodeWeavers CrossOver Office for, 241, 258–259
 - finding Windows-equivalents, 257
 - VMware for, 261–262
 - Wine for, 262–264
 - Win4Lin program for, 259–260
 - Windows 2000
 - determining hard drive free space, 27–28
 - finding hardware information, 36–37
 - viewing partitions, 23–24
 - Windows Vista
 - determining hard drive free space, 27
 - finding hardware information, 38
 - partitioning software for, 20
 - viewing partitions, 22–23
 - Windows XP
 - determining hard drive free space, 27–28
 - finding hardware information, 36–37
 - `fixmbr` command, 64
 - mapping Samba print server, 346
 - removing Linux with, 64
 - viewing partitions, 23–24
 - Windows-to-Linux connections with SSH, 308–310
 - Wine software, 262–264
 - Win4Lin program, 20, 259–260
 - WinModems, 38–39, 137
 - WinZip program, 269
 - wireless networking. *See also* WAP (wireless access point)
 - ad-hoc mode, 318
 - bridge for, 139, 325–329
 - configuring remote devices, 331–333
 - further information, 138–139, 148
 - installing your wireless NIC, 321–322
 - Internet connection challenges, 138, 148
 - Internet connection setup, 148–151
 - managed mode, 318
 - ndiswrapper for, 148, 322–323
 - NIC configuration, 149–151, 320–321
 - planning your network, 319
 - router for, 139
 - saving bridging commands, 328–329
 - security configuration, 330–331
 - testing, 329–330, 333
 - troubleshooting, 333–334
 - WMV (Windows Media Video) format, 240–241, 244
 - word processors
 - in office suites, 191
 - text editors versus, 188
 - Writer (OpenOffice.org), 71, 193–197
 - write (`w`) permission, 109
 - Writer (OpenOffice.org)
 - documents in, 196
 - formats supported, 196–197
 - formatting toolbar, 196
 - GNOME icon for, 71
 - menu bar, 193–194
 - quick-launch icon for, 192
 - ruler, 196
 - standard toolbar, 194–195
- ✕ •
- × (execute) permission, 109, 110
 - Xandros distribution, 16–17, 64
 - X-Chat IRC program, 184
 - xine video player, 242
 - X.org.0.log file, 313
 - XP. *See* Windows XP
- ♪ •
- Young, Robert (Red Hat cofounder), 16
 - Yum software installer
 - for Firefox plug-ins, 241
 - installing xine video player using, 242
 - nightly automatic updates using, 305
 - for 3-D video card drivers, 254